arXiv:1602.01295v1 [cs.DS] 3 Feb 2016

HOW PROOFS ARE PREPARED AT CAMELOT*

ANDREAS BJÖRKLUND AND PETTERI KASKI

Abstract. We study a design framework for robust, independently verifiable, and workload-balanced distributed algorithms working on a common input. An algorithm based on the framework is essentially a *distributed encoding procedure* for a Reed–Solomon code (each compute node is tasked to produce one or more symbols of the codeword), which enables (a) robustness against byzantine adversarial failures with intrinsic error-correction and identification of failed nodes, and (b) independent randomized verification to check the entire computation for correctness, which takes essentially no more resources than each node individually contributes to execute the computation. The framework also enables smooth tradeoffs between per-node compute time and the number of nodes used for computation.

The framework builds on recent *Merlin–Arthur* proofs of batch evaluation of Williams [*Electron. Colloq. Comput. Complexity*, Report TR16-002, January 2016] with the basic observation that *Merlin's magic is not needed* for batch evaluation—mere *Knights* can prepare the independently verifiable *proof*, in parallel, and with intrinsic error-correction. Dually, *should Merlin materialize*, he can relieve the Knights and instantaneously supply the proof, in which case these algorithms are, *as is*, Merlin–Arthur protocols.

The contribution of this paper is to show that in many cases the verifiable batch evaluation framework admits algorithms that match in total resource consumption the best known sequential algorithm for *solving* the problem. As our main technical result, we show that the k-cliques in an n-vertex graph can be counted *and* verified in per-node $O(n^{(\omega+\epsilon)k/6})$ time and space on $O(n^{(\omega+\epsilon)k/6})$ compute nodes, for any constant $\epsilon > 0$ and positive integer k divisible by 6, where $2 \leq \omega < 2.3728639$ is the exponent of square matrix multiplication over the integers. This matches in total running time the best known sequential algorithm, due to Nešetřil and Poljak [*Comment. Math. Univ. Carolin.* 26 (1985) 415–419], and considerably improves its space usage and parallelizability. Further results include novel algorithms for counting triangles in sparse graphs, computing the chromatic polynomial of a graph, and computing the Tutte polynomial of a graph.

*The research leading to these results has received funding from the Swedish Research Council grant VR 2012-4730 "Exact Exponential-Time Algorithms" (A.B.) and the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement 338077 "Theory and Practice of Advanced Search and Enumeration" (P.K.). Work done in part while the authors were visiting the Simons Institute for the Theory of Computing.

1. Introduction

1.1. **A scene of distress and relief at Camelot.** Picture K equally capable Knights seated around the Round Table, distressed. At the center of the table stands the Input. The Knights have been tasked to prepare a Proof about the virtues of the input, and to make *extreme* haste: a proof must be prepared in time T/K, where T is the *fastest time known in all of Britannia* for any single soul to *reveal* such subtle virtues, let alone give *uncontestable proof* thereof.

Nigh impossible is the task in fact, for the Lady Morgana has enchanted many a poor Knight with her cunning dark magic, yet virtues must be revealed and proof thereof prepared. And not just any proof, but one that any lone soul can check, in time T/K and with a few tosses of a fair coin, and next to never be convinced if foul virtues it claims.

However, not all is lost, for the Knights recall the teachings of the wizard Merlin[1] and his powerful proofs that King Arthur so treasures. So the Knights agree to together evaluate in batch, and to individually decode-and-check for signs of the darkest of magic...

1.2. **A template for community computation over common input.** Building on Merlin–Arthur proofs [5] of *batch evaluation* of Williams [35], this paper documents a "Camelot template" for verifiable distributed computation that is robust against *adversarial byzantine failures* at the nodes and produces a static, independently verifiable *proof* that the computation succeeded. In essence, we observe that the act of *preparing* a proof of batch evaluation in the Williams framework is *magicless*[2] and *robust* against errors. Our contribution is in particular to show that for a number of problems such robust and verifiable algorithms essentially match in total resource consumption the best known sequential algorithm for *solving* the problem.

By robustness we mean here that during proof preparation a substantial number[3] of nodes may undergo byzantine failure and a correct proof will still be produced at each correctly functioning node. (We refer to e.g. Tiwari *et al.* [31] for recent empirical motivation why robustness is an increasingly desirable goal in algorithm design.)

By verifiability we mean here that if the computation *did* fail, each node detects this individually *regardless of how many nodes experienced byzantine failure* during computation. In more precise terms, each node (or other entity) can verify the static proof against the input so that the probability

[1] Alas, the good wizard is taking a vacation, hence the distress at Camelot.

[2] In Arthurian parlance, *Merlin's magic is not needed—mere Knights* prepare the *proof*, in parallel, and with intrinsic error-correction. Dually, *should Merlin materialize*, he can relieve the Knights and instantaneously supply the proof, in which case these algorithms are, *as is*, Merlin–Arthur protocols.

[3] In precise terms, up to the decoding limit of the decoder used for the Reed–Solomon code, such as at most $(e - d - 1)/2$ nodes if we use, e.g., the Gao [17] decoder and e evaluation points for a proof polynomial of degree d; see §2.3.

of accepting an incorrect proof is negligible *and* controllable by the verifier *after* the proof is presented for verification. Furthermore, the resources needed for verification are essentially identical to the resources that a single node contributes to the community.

1.3. **How Camelot prepares, corrects, and checks a proof.** Let us interpret the Williams batch evaluation framework [35] as a template for distributed computing with K compute nodes. The key idea is to task the K nodes to make *evaluations* of a univariate polynomial

$$P(x) = p_0 + p_1 x + p_2 x^2 + \ldots + p_d x^d \pmod{q},$$

where q is a prime that we assume each node can easily compute by looking at the common input.[4] (We also assume that each node can easily compute an upper bound for d from the common input.)

The polynomial $P(x)$ has been carefully constructed so that (a) it reveals[5] the desired properties of the input *and* (b) the *coefficients* $p_0, p_1, \ldots, p_d \in \mathbb{Z}_q$ constitute a probabilistically verifiable *proof* that these properties have been correctly computed. In fact, the latter property is *intrinsic* to the template.

1. Proof preparation, in distributed encoded form. We assume that the input (the problem instance) is available at each of the K nodes; for example, the input may be broadcast over a network to all the nodes. For a prime q and an integer e with $d + 1 \leq e \leq q$, let us collectively task the K nodes to compute the sequence of e evaluations

$$P(0), P(1), \ldots, P(e-1) \pmod{q} \tag{1}$$

so that each node is responsible for about e/K evaluations, and broadcasts its evaluations to the other nodes once they are ready.[6] This sequence of evaluations is, *by definition*, an encoding of the proof in the classical nonsystematic *Reed–Solomon code* [28] (see §2). From the perspective of *community computation*, this is a serendipitous property:

2. Error-correction during preparation of the proof. It is possible to error-correct a proof *in preparation* to cope with failing nodes in the community of nodes, up to the error-correction capability of the code and the decoding algorithm. Indeed, since the nodes are tasked to produce entries in a Reed–Solomon codeword, each node can *individually* run a *fast* Reed–Solomon decoder [17] (see §2) on the entries it has received from the community

[4]For ease of exposition, we work over fields of prime order though generalizations to field extensions are possible, e.g., to obtain better fault tolerance. Further dimensions for generalizations include, e.g., the use of multivariate (Reed–Muller [24, 27]) polynomial codes.

[5]Possibly requiring iteration and postprocessing, e.g., the use of multiple distinct primes q and the Chinese Remainder Theorem to reconstruct large integers.

[6]In particular, in total e integers modulo q are broadcast over the network, with the objective that all nodes receive the e evaluations. (If only a single node is to error-correct and verify the proof, then no broadcasting is necessary in this step – each node forwards its evaluations only to the verifier node.)

to recover *both* the actual proof $p_0, p_1, \ldots, p_d \in \mathbb{Z}_q$ *and* the evaluations $P(0), P(1), \ldots, P(e-1) \in \mathbb{Z}_q$ *that failed*, that is, the *error locations* identified by the decoder. Thus, if a node so desires, it can, on its own, identify the nodes that did not properly participate in the community effort by looking at the error locations. Furthermore,[7] all nodes can reach the same *collective* conclusion about the community contribution on their own.

3. Checking a putative proof for correctness. Each node (or any other entity) with access to the common input and to a sequence $\tilde{p}_0, \tilde{p}_1, \ldots, \tilde{p}_d \in \mathbb{Z}_q$ can, probabilistically, check whether the sequence is in fact the correct proof $p_0, p_1, \ldots, p_d \in \mathbb{Z}_q$. All a node has to do is to select a uniform random integer $x_0 \in \mathbb{Z}_q$ and accept the proof if and only if

$$P(x_0) = \sum_{j=0}^{d} \tilde{p}_j x_0^j \pmod{q} \tag{2}$$

holds, at the cost of one evaluation of $P(x)$. (That is, the verifier executes the same algorithm that the nodes use for preparing the proof.[8]) If the sequence $\tilde{p}_0, \tilde{p}_1, \ldots, \tilde{p}_d$ *is* the correct proof, the verifier always accepts. If the sequence is incorrect, it follows immediately from the fundamental theorem of algebra that $P(x_0) \neq \sum_{j=0}^{d} \tilde{p}_j x_0^j \pmod{q}$ holds with probability at least $1 - d/q$. Furthermore, since the verifier knows q and an upper bound for d, the verifier can control the probability of accepting an incorrect proof by independent repetitions of the test.

1.4. **Optimality relative to best known sequential algorithms.** Ideally, a Camelot algorithm should match in *total*[9] running time the worst-case running time of the best known *sequential* algorithm for *solving* the problem. By design, such a Camelot algorithm not only solves the problem, but does so in *parallel*, *robustly*, and gives *an independently verifiable proof* that it succeeded. What is more, the parallel part of the computation (1) consists of evaluating *the same* polynomial P at distinct points, making the parallel computations intrinsically *workload-balanced* and ideal for *vector-parallelization* (*single-instruction multiple-data (SIMD) parallelization*) since each evaluation follows the same instruction stream.

In essence, we seek the following *optimal speedup tradeoff* for a Camelot design. Assume that the best known sequential algorithm solves the problem

[7] In precise terms, a collective conclusion can be reached *assuming the nodes agree on the evaluations* $P(0), P(1), \ldots, P(e-1) \in \mathbb{Z}_q$ that were produced by the community. It is perhaps a realistic practical assumption that such an agreement can be reached by an appropriate broadcast mechanism for informing all the nodes about a node's contributions. We observe that such collective agreement *need, however, not be available for purposes of decoding the proof*—the decoder works as long as it receives a sufficient number of correct entries of the codeword, and these entries may be different for different nodes.

[8] In precise terms, this is to obtain the left-hand side of (2). The right-hand side can be computed from $\tilde{p}_0, \tilde{p}_1, \ldots, \tilde{p}_d \in \mathbb{Z}_q$ using Horner's rule.

[9] The sum of running times of all the nodes.

essentially in time T. Then, an optimal Camelot design for K nodes should run in $E = T/K$ proof construction time and space for each node, and hence E evaluation (proof verification) time for a proof of size K, for any $K \leq T^{1/2}$. Here the upper bound $T^{1/2}$ is intrinsic to the framework since each node receives the whole proof for individual decoding and hence $E \geq K$.[10]

In practice, our subsequent algorithm designs will not attain the optimal tradeoff for all problems considered. *However, for a number of canonical problems we are able to attain the optimal tradeoff.*

1.5. **Our results—four highlights.** This section summarizes the four main technical results of this paper that highlight the technical versatility of the Camelot framework. For ease of exposition we present our results as non-tradeoff versions with the maximum value of K (and hence the fastest wall-clock runtime E to prepare and check the proof) that we can attain, with the understanding that a smooth tradeoff up to the maximum K is possible and can be read off the proofs of our results.

When stating our results, we refer to K as the *proof size* and E as the *(running) time* of a Camelot algorithm. Put otherwise, the time E of a Camelot algorithm is the *wall-clock* time to prepare and verify the proof using K compute nodes working in parallel on the common input. The *total time* used by all the nodes is EK. The underlying precise degree d of the proof polynomial $P(x)$ is, up to factors suppressed by asymptotic notation, identically bounded as the stated proof size.[11] With asymptotic notation we follow the convention that $\tilde{O}(\cdot)$ suppresses a factor polylogarithmic in the input size and $O^*(\cdot)$ suppresses a factor polynomial in the input size.

Each Camelot algorithm defines, *as is*, a *Merlin–Arthur* protocol [5], where Merlin's proof has the stated size, and Arthur's verification time is the stated time for the Camelot algorithm.

1. Optimal Camelot algorithms in polynomial time: counting small cliques. Our main result is that the problem of counting all the k-cliques in a graph admits a Camelot algorithm *that attains the optimal tradeoff relative to the best known sequential algorithm.* Let us write $2 \leq \omega < 2.3728639$ for the exponent of square matrix multiplication over the integers [23].

Theorem 1. *For any constant $\epsilon > 0$ and any positive integer k divisible by 6, the number of k-cliques in an n-vertex graph can be computed with a Camelot algorithm that constructs a proof of size $O\big(n^{(\omega+\epsilon)k/6}\big)$ in time $O\big(n^{(\omega+\epsilon)k/6}\big)$.*

[10] This of course does not rule out other frameworks where a single node does not work with the entire proof.

[11] The parameters e and q can be chosen with similar asymptotic bounds as d. We choose to omit the details of these parameters to keep the statement of our results concise, please consult the proofs for precise bounds on d to enable fine-tuning of selection of $d + 1 \leq e \leq q$ as appropriate for desired error-tolerance (during proof preparation) and soundness guarantee (for proof verification).

The best known sequential algorithm, due to Nešetřil and Poljak [25] (see also Eisenbrand and Grandoni [12]), runs in time $O(n^{(\omega+\epsilon)k/3})$ and space $O(n^{2k/3})$. In particular, Theorem 1 gives an optimal tradeoff and Theorem 2 improves the space usage and parallelizability of the best known sequential algorithm.

The algorithm in Theorem 1 is based on a novel, more space-efficient family of arithmetic circuits for evaluating a particular $\binom{6}{2}$-linear form that enables a very space-efficient design for counting 6-cliques in a graph. We state the latter result independently of the Camelot framework:

Theorem 2. *For any constant $\epsilon > 0$ and any positive integer k divisible by 6, there exists an algorithm that counts the k-cliques in a given n-vertex graph in time $O\big(n^{(\omega+\epsilon)k/6}\big)$ and space $O\big(n^{k/3}\big)$. Furthermore, the algorithm can be executed in per-node time $O\big(n^{(\omega+\epsilon)k/6}\big)$ and space $O\big(n^{k/3}\big)$ on $O\big(n^{(\omega+\epsilon)k/6}\big)$ parallel compute nodes.*

Observe in particular that for $k = 6$ the space usage in Theorem 2 is $O(n^2)$, that is, the algorithm uses asymptotically no more space than the input graph.

2. Sparsity-aware Camelot algorithms: counting triangles. As our second highlight we show that Camelot algorithms can be made aware of *sparsity* in the given input. In particular, we study the task of counting triangles in a given n-vertex graph. Taking the $O(n^{\omega+\epsilon})$-time sequential algorithm of Itai and Rodeh [19] as our baseline for parallelization, we present a Camelot algorithm whose running time is essentially linear in the size of the input.

Theorem 3. *For any constant $\epsilon > 0$, the number of triangles in an n-vertex graph with $m \geq n^{(\omega+\epsilon)/2}$ edges can be computed with a Camelot algorithm that constructs a proof of size $O\big(n^{\omega+\epsilon}/m\big)$ in time $\tilde{O}\big(m\big)$.*

The algorithm in Theorem 3 is based on a novel, more space-efficient split/sparse polynomial extension of Yates's algorithm [36] for computing matrix–vector products when the matrix has iterated Kronecker product structure. We state a space-efficient and parallel version of the result independently of the Camelot framework:

Theorem 4. *For any constant $\epsilon > 0$, there exists an algorithm that counts the triangles in a given connected n-vertex graph with m edges in time $O\big(n^{\omega+\epsilon}\big)$ and space $\tilde{O}\big(m\big)$. Furthermore, the algorithm can be executed in per-node time and space $\tilde{O}\big(m\big)$ on $O\big(n^{\omega+\epsilon}/m\big)$ parallel compute nodes.*

Remark. For a *sparse* input consisting of m edges, the fastest known sequential algorithm for triangle counting due to Alon, Yuster, and Zwick [4] runs in time $O\big(m^{2(\omega+\epsilon)/(\omega+1)}\big)$. Our designs in Theorem 3 and Theorem 4 only match the weaker Itai–Rodeh bound $O(n^{\omega+\epsilon})$; matching the Alon–Yuster–Zwick bound with an optimal Camelot algorithm remains an open problem.

If we seek only a parallel algorithm, the split/sparse variant of Yates's algorithm enables us to match the Alon–Yuster–Zwick bound:

Theorem 5. *For any constant $\epsilon > 0$, there exists an algorithm that counts the triangles in a given graph with m edges in time $O\big(m^{2(\omega+\epsilon)/(\omega+1)}\big)$ and space $\tilde{O}\big(m\big)$. Furthermore, the algorithm can be executed in per-node time and space $\tilde{O}\big(m\big)$ on $O\big(m^{(\omega-1+\epsilon)/(\omega+1)}\big)$ parallel compute nodes.*

3. Optimal Camelot algorithms in exponential time: the chromatic polynomial. As our third highlight result, we show that the problem of computing the chromatic polynomial of a given graph admits an optimal Camelot algorithm.

Theorem 6. *The chromatic polynomial of an n-vertex graph can be computed with a Camelot algorithm that constructs a proof of size $O^*(2^{n/2})$ in time $O^*(2^{n/2})$.*

Here the best known sequential algorithm runs in $O^*(2^n)$ time and $O^*(1.292^n)$ space [7, 9]. In particular, Theorem 6 gives an optimal tradeoff.

4. Near-optimal Camelot algorithms in exponential time: the Tutte polynomial. The *Tutte polynomial* or the *bichromatic polynomial* of a graph is a universal invariant for all constant-coefficient deletion–contraction recurrences. The Tutte polynomial subsumes a large number of #P-hard counting problems, such as the chromatic polynomial, the flow polynomial, the all-terminal reliability polynomial, the partition functions of the Ising and Potts models in statistical physics, the Jones polynomial of a link in knot theory, and more. (We refer to Welsh [33] and Godsil and Royle [15] for a detailed discussion on polynomial graph invariants and universality.)

For the Tutte polynomial, we can almost attain the optimal tradeoff:

Theorem 7. *For any constant $\epsilon > 0$, the Tutte polynomial of an n-vertex graph can be computed with a Camelot algorithm that constructs a proof of size $O^*(2^{n/3})$ in time $O^*(2^{(\omega+\epsilon)n/3})$ and per-node space $O^*(2^{2n/3})$.*

Here the main bottleneck is the *space usage* of the best known sequential algorithm [8], which we must reduce from $O^*(2^n)$ to $O^*(2^{2n/3})$ to arrive at a design that can be placed in the Camelot framework. In the process we rely on a tripartite decomposition of Williams (originally presented in the context of 2-CSPs [34]) to decompose an edge-enumerator function, which forces a dependence on fast matrix multiplication in the per-node running times. The best known sequential algorithm runs in $O^*(2^n)$ time and space [8]. Thus, up to polynomial factors in n, the design in Theorem 7 attains the optimal tradeoff $T = KE$ only if $\omega = 2$. Furthermore, the design enables parallelization up to K nodes for $K \leq T^{1/3}$ only.

1.6. **Overview of techniques, earlier work, and further results.** The present Camelot framework can be seen as a continuation of two recent

works aimed at understanding the fine-grained[12] complexity of computational problems and proof systems, initiated in a nondeterministic setting with a deterministic proof verifier by Carmosino *et al.* [11] and with the subsequent breakthrough by Williams [35] establishing nontrivial qualitative bounds for a *randomized verifier*. From this perspective, the Camelot framework pursues an even more fine-grained goal by focusing also on the effort of the *prover*, with the objective of relating the *total effort to prepare the proof* to the best known sequential algorithms for *solving*[13] the problem at hand.

To design a Camelot algorithm, all it takes is to come up with the proof polynomial P and a fast evaluation algorithm for P. Our main results outlined in the previous section demonstrate designs that attain either optimal or near-optimal speedup compared with the best known sequential algorithms, where the sequential algorithms themselves are rather technical.

Many further designs are possible, and we find it convenient to introduce the techniques underlying our highlight results gradually, starting from more accessible designs, and with a further motivation to illustrate the generality of the Camelot framework.

Perhaps the more accessible Camelot designs can essentially be read off from the Merlin–Arthur proofs of batch evaluation of Williams [35]. We condense three such designs into the following theorem. In each case the problem is #P-hard [32], and no algorithm with running time $O^*(c^n)$ for $c < 2$ is known; algorithms with sub-exponential speedup over $O(2^n)$ are known, however [1, 6].

Theorem 8. *There exist Camelot algorithms that construct a proof of size* $O^*(2^{n/2})$ *in time* $O^*(2^{n/2})$ *for each of the following:*

(1) *The number of solutions of a CNF formula on* n *variables.*
(2) *The permanent of an* $n \times n$ *integer matrix.*
(3) *The number of Hamilton cycles in an* n*-vertex graph.*

(We prove Theorem 8 in Appendix A.)

The proof polynomials P associated with Theorem 8 follow a design pattern where (i) the desired counts can be computed in essentially linear time from the available evaluations of P (assuming the evaluations are correct!), and (ii) the polynomial P is a composition of a vector of univariate interpolating polynomials and a multivariate polynomial.

Armed with this insight from batch evaluation, further polynomials and algorithms can be extracted from earlier work on algebraic algorithms, such as the partitioning and covering framework of Björklund *et al.* [7]. The

[12] *Fine-grained* in the sense that the objective is to obtain more precise *quantitative* bounds on resource usage compared with the more traditional coarse-grained and qualitative division between, e.g., polynomial time and exponential time in the context of classical Merlin–Arthur proofs [5].

[13] Without the burden of preparing a proof of correctness.

following result with an optimal tradeoff can be read off from the inclusion–exclusion formulas in this framework.

Theorem 9. *The number of t-element set covers constructible from a given family of $O^*(1)$ subsets of an n-element ground set can be computed with a Camelot algorithm that constructs a proof of size $O^*(2^{n/2})$ in time $O^*(2^{n/2})$.*

(We prove Theorem 9 in Appendix A.)

Our Camelot algorithms for the chromatic polynomial (Theorem 6) and the Tutte polynomial (Theorem 7) require a more subtle construction for the proof polynomial P to enable fast distributed evaluation at the nodes. In essence this is because the evaluation algorithm at each node must work with a structured set family that may have size up to 2^n (for example, the family of independent sets of a graph), requiring that we rely on the structure in the family to arrive at $O^*(2^{n/2})$ evaluation time. Here our design is guided by the parallel algorithm of Björklund *et al.* [9] for computing the chromatic polynomial. Also the structure of the proof polynomial P will be rather different from the polynomials underlying Theorem 8 in two respects: (a) the property that we want to compute is one *coefficient* of P, and (b) to obtain a compact univariate proof, we will rely on Kronecker substitution and weight tracking.

To present the design underlying Theorems 6 and 7 in more technically accessible parts, as a warmup it will be convenient to first establish an exact-covering version of Theorem 9 that takes in an *exponential-sized* family of subsets.

Theorem 10. *The number of t-element set partitions constructible from a given family of $O^*(2^{n/2})$ subsets of an n-element ground set can be computed with a Camelot algorithm that constructs a proof of size $O^*(2^{n/2})$ in time $O^*(2^{n/2})$.*

The proofs of Theorems 10, 6, and 7 all rely on the same template for the proof polynomial P which we introduce in §7, followed by the proofs in §8, §9, and §10, respectively.

Let us now move from #P-hard problems to problems in polynomial time. Again it is convenient to start with a number of more accessible designs followed by our main result. The first two designs can be read off the batch evaluation results of Williams [35]; the third design is a minor extension of the same ideas. In each case no sequential algorithm with running time $O(n^{2-\epsilon}t^c)$ for positive constants ϵ and c is known; sub-polynomial speedups in n are known [1, 3, 18, 26].

Theorem 11. *There exist Camelot algorithms that construct a proof of size $\tilde{O}(nt^c)$ in time $\tilde{O}(nt^c)$ for each of the following:*

(1) *The number of pairs of orthogonal vectors in a given set of n Boolean vectors of dimension t. (With $c = 1$.)*

(2) *The distribution of Hamming distances between two given sets of n Boolean vectors of dimension t. (With $c = 2$.)*
(3) *The number of solutions of an instance of Convolution3SUM consisting of n integers with bit length t. (With $c = 2$.)*

(We prove Theorem 11 in Appendix A.)

The proof polynomials underlying Theorem 11 again consist of interpolation polynomials composed with a multivariate polynomial designed to indicate combinations of interest, where some care is required to control the constants c. Designs analogous to Theorem 11 lead to Camelot algorithms for counting k-cliques in an n-vertex graph that have proof size $\tilde{O}(n^{k/2})$ and run in time $n^{k/2+O(1)}$. (See Williams [35].)

In contrast, our design for Theorem 1 relies on a novel arithmetic circuit design for counting 6-cliques that nests *four* levels of fast matrix multiplication so that the products at the top-level multiplication can (a) be evaluated in parallel *and* (b) extended to a proof polynomial P by careful use of Yates's algorithm [36] and fast matrix multiplication to enable efficient evaluation. The top-level multiplications rely on a *trilinear* representation of the matrix multiplication tensor $\langle n, n, n\rangle$ to assemble the monomials of a $\binom{6}{2}$-linear form out of a 3-linear base and three 4-linear parts, whose consistency is obtained via the trilinear representation. In essence, we show that 6-cliques can be counted with an arithmetic circuit of essentially the *same size* but with considerably more *depth* than the Nešetřil–Poljak [25] circuit.

We present the proof of Theorem 1 in two parts. First, in §4 we present the novel arithmetic circuits for the $\binom{6}{2}$-linear form. Then, in §5 we first prove Theorem 2 as a warmup and then extend the top-level products in the circuit design to a proof polynomial P and its fast evaluation algorithm.

The arithmetic circuits and the proof polynomial for the $\binom{6}{2}$-linear form admit also further applications beyond counting subgraphs, including further applications with #P-hard problems. For example, as a corollary of §4 and §5 we obtain an optimal Camelot algorithm for enumerating variable assignments by the number of satisfied constraints in an instance of 2-constraint satisfaction. The trivial sequential algorithm runs in time $O^*(\sigma^n)$ for n variables that take values over an alphabet of σ symbols, and the best known sequential algorithm runs in time $O^*(\sigma^{(\omega+\epsilon)n/3})$ for any constant $\epsilon > 0$ [34].

Theorem 12. *For any constant $\epsilon > 0$ there exists a Camelot algorithm that constructs a proof of size $O^*(\sigma^{(\omega+\epsilon)n/6})$ in time $O^*(\sigma^{(\omega+\epsilon)n/6})$ for the enumeration of variable assignments by the number of satisfied constraints for a given system of 2-constraints over n variables, with each variable taking values over an alphabet of size σ.*

(We prove Theorem 12 in Appendix B.)

Remark. Theorem 12 admits a generalization to weighted instances where each 2-constraint has a nonnegative integer weight at most W. In this case both the proof size and the per-node running time get multiplied by W.

Let us now discuss the techniques for working with sparsity and fast matrix multiplication underlying Theorem 3 and Theorem 4. As with Theorem 1 and Theorem 2, the technical contribution is to split a trilinear form of matrix multiplication into independent parts that can be executed in parallel, so that each part takes advantage of the sparse input. Here the key technical tool is a novel split/sparse variant of Yates's algorithm that we will present in §3.2. This algorithm applied to the Itai–Rodeh design [19] gives Theorem 4 as an essentially immediate corollary. Furthermore, the split/sparse algorithm admits a *polynomial extension* (presented in §3.3) which in turn enables the proof polynomial and the Camelot algorithm in Theorem 3.

Let us conclude this section by observing that all the sequential algorithms that we use as baseline designs to our Camelot designs are deterministic algorithms. However, the Camelot framework extends in a natural way to randomized algorithms and proving the correctness of a randomized computation if we assume the nodes have access to a public random string. In particular, the computation for any outcome of the random string is deterministic and hence verifiable in the deterministic framework.

2. Preliminaries

This section recalls notation and background results used in our main development.

2.1. **Notation and conventions.** For a logical proposition Ψ we use Iverson's bracket notation $[\Psi]$ to indicate a 1 if Ψ is true and a 0 if Ψ is false. At the risk of slight ambiguity in notation, for a nonnegative integer n we will, however, write $[n]$ to indicate the set $\{1, 2, \ldots, n\}$. The notation $\tilde{O}(\cdot)$ hides a factor polylogarithmic in the input size. The notation $O^*(\cdot)$ hides a factor polynomial in the input size. The product over an empty set evaluates to 1.

2.2. **Fast polynomial arithmetic.** Let us recall the fast arithmetic toolbox (see von zur Gathen and Gerhard [16]) for polynomials with coefficients in a finite field $\mathbb{F}_q$. Computational complexity is measured in the number of arithmetic operations in $\mathbb{F}_q$.

Multiplication and division of two polynomials of degree at most d can be computed in $O(d \log d \log\log d)$ operations. The greatest common divisor of two polynomials of degree at most d can be computed in $O(d \log^2 d \log\log d)$ operations.[14] Given a polynomial

$$P(x) = p_0 + p_1 x + p_2 x^2 \ldots + p_d x^d \in \mathbb{F}_q[x]$$

[14] At the same cost one can obtain any single remainder in the sequence of remainders of the extended Euclidean algorithm, together with the multiplier polynomials that produce the remainder. This is serendipitous for the Gao [17] decoder which we review in §2.3.

and points $x_0, x_1, \ldots, x_d \in \mathbb{F}_q$ as input, the *evaluation map*

$$(P(x_0), P(x_1), \ldots, P(x_d)) \in \mathbb{F}_q^{d+1}$$

can be computed in $O(d \log^2 d \log\log d)$ operations. Dually, given distinct points $x_0, x_1, \ldots, x_d \in \mathbb{F}_q$ and values $y_0, y_1, \ldots, y_d \in \mathbb{F}_q$ as input, the *interpolation map* computes the coefficients $(p_0, p_1, \ldots, p_d) \in \mathbb{F}_q^{d+1}$ of a polynomial $P(x)$ such that $P(x_i) = y_i$ for all $i = 0, 1, \ldots, d$ in $O(d \log^2 d \log\log d)$ operations.

2.3. **Reed–Solomon codes, fast encoding and decoding.** We work with univariate polynomial codes in nonsystematic form as originally discovered by Reed and Solomon [28] to allow for fast decoding based on polynomial arithmetic (cf. Gao [17]).

Fix a finite field $\mathbb{F}_q$ and any $1 \leq e \leq q$ distinct elements $x_1, x_2, \ldots, x_e \in \mathbb{F}_q$. The parameter e is the *length* of the code. To *encode* a *message* $(p_0, p_1, \ldots, p_d) \in \mathbb{F}_q^{d+1}$ consisting of $d+1$ *symbols*, $1 \leq d+1 \leq e$, form the message polynomial $P(x) = p_0 + p_1 x + p_2 x^2 + \ldots + p_d x^d$ and the *codeword*

$$(P(x_1), P(x_2), \ldots, P(x_e)) \in \mathbb{F}_q^e .$$

The codeword can be computed from the message in $O(e \log^2 e \log\log e)$ operations using fast interpolation (see §2.2).

Since a nonzero polynomial of degree d has at most d roots, any two codewords either agree in at most d symbols or are identical. Thus, we can uniquely *decode* the message from any *received* word $(r_1, r_2, \ldots, r_e) \in \mathbb{F}_q^n$ that differs from the codeword in at most $(e-d-1)/2$ symbols.

To decode the received word (or assert decoding failure), we use the following algorithm of Gao [17]. As a precomputation step, compute the polynomial $G_0(x) = \prod_{i=1}^{e}(x - x_i) \in \mathbb{F}_q[x]$.[15] To decode $(r_1, r_2, \ldots, r_e) \in \mathbb{F}_q^e$, first interpolate the unique polynomial $G_1(x) \in \mathbb{F}_q[x]$ of degree at most $e-1$ with $G_1(x_i) = r_i$ for all $i = 1, 2, \ldots, e$. Second, apply the extended Euclidean algorithm to the polynomials $G_0(x)$ and $G_1(x)$, but stop as soon as the remainder $G(x)$ has degree less than $(e+d+1)/2$. At this point suppose that

$$U(x)G_0(x) + V(x)G_1(x) = G(x) .$$

Finally, divide $G(x)$ by $V(x)$ to obtain

$$G(x) = P(x)V(x) + R(x) ,$$

where the degree of $R(x)$ is less than the degree of $V(x)$. If $R(x) = 0$ and $P(x)$ has degree at most d, output $P(x)$ as the result of decoding; otherwise assert decoding failure. We observe that decoding runs in $O(e \log^2 e \log\log e)$ operations using fast polynomial arithmetic.

[15] When $e = q$, we have $G_0(x) = x^q - x$ which suffices for our applications in this paper.

3. A split/sparse variant of Yates's algorithm

This section presents a variant of Yates's algorithm [36] that accepts *sparse* input and *splits* its output into multiple parts that can be produced independently of each other, in parallel. This variant can essentially be read off the split/sparse fast zeta transform of Björklund *et al.* [9]. Furthermore, the algorithm readily admits a polynomial extension that will prove serendipitous for the design of Camelot algorithms (see §6).

3.1. **Yates's algorithm.** Let us start by recalling the classical Yates's algorithm. For positive integers s, t, k, Yates's algorithm [36] multiplies a given $s^k \times 1$ vector x with a structured $t^k \times s^k$ matrix M to produce as output a $t^k \times 1$ vector $y = Mx$. The structural assumption about the matrix M is that it is a Kronecker power $M = A^{\otimes k}$ of a small $t \times s$ matrix A with entries α_{ij} for $i = 1, 2, \ldots, t$ and $j = 1, 2, \ldots, s$. All arithmetic takes place in a field $\mathbb{F}$.

For positive integers d, k we shall tacitly identify each $j \in [d^k]$ with $(j_1, j_2, \ldots, j_k) \in [d]^k$. Indeed, we can view j as a k-digit integer in base d, where $j_1, j_2, \ldots, j_k$ are the k digits.

Given x and A as input, Yates's algorithm seeks to compute, for all $i \in [t^k]$, the value

$$y_i = \sum_{j \in [s^k]} \alpha_{i_1 j_1} \alpha_{i_2 j_2} \cdots \alpha_{i_k j_k} x_j \,. \tag{3}$$

The algorithm proceeds in k levels, where level $\ell = 1, 2, \ldots, k$ takes as input a $t^{\ell-1} s^{k+1-\ell} \times 1$ vector $y^{(\ell-1)}$ and produces as output a $t^\ell s^{k-\ell} \times 1$ vector $y^{(\ell)}$. The input to the first level is $y^{(0)} = x$. The vector $y^{(\ell)}$ is defined for all $i_1, i_2, \ldots, i_\ell \in [t]$ and $j_{\ell+1} j_{\ell+2} \cdots j_k \in [s]$ by

$$y^{(\ell)}_{i_1 i_2 \cdots i_\ell j_{\ell+1} j_{\ell+2} \cdots j_k} = \sum_{j_1, j_2, \ldots, j_\ell \in [s]} \alpha_{i_1 j_1} \alpha_{i_2 j_2} \cdots \alpha_{i_\ell j_\ell} x_j \,. \tag{4}$$

It is immediate that $y^{(k)} = y$ as desired. To compute $y^{(\ell)}$ given $y^{(\ell-1)}$ as input, we observe that for all $i_1, i_2, \ldots, i_\ell \in [t]$ and $j_{\ell+1} j_{\ell+2} \cdots j_k \in [s]$ it holds that

$$y^{(\ell)}_{i_1 i_2 \cdots i_\ell j_{\ell+1} j_{\ell+2} \cdots j_k} = \sum_{j_\ell \in [s]} \alpha_{i_\ell j_\ell} y^{(\ell-1)}_{i_1 i_2 \cdots i_{\ell-1} j_\ell j_{\ell+1} \cdots j_k} \,. \tag{5}$$

In essence, (5) proceeds from the input $x = y^{(0)}$ towards the output $y = y^{(k)}$ by evaluating "one of the k nested sums at a time", as is readily verified using induction on ℓ and (4).

From (4) it follows that Yates's algorithm runs in $O((s^{k+1} + t^{k+1})k)$ operations and working space for $O(s^k + t^k)$ field elements.

3.2. **The split/sparse algorithm.** Let us assume that $t \geq s$. Suppose the input vector x is *sparse* so that $x_j \neq 0$ only if $j \in D$ for a nonempty set $D \subseteq [s^k]$. Accordingly, we assume that the input consists of $O(|D|)$ field

elements. We seek to produce the t^k elements of the output vector y in roughly $t^k/|D|$ independent parts so that each part consists of roughly $|D|$ entries of y. In more precise terms, the number of entries of y in each part is the minimum multiple of t at least $|D|$.

The algorithm in particular *splits* the output indices $i_1 i_2 \cdots i_k$ into two groups, $i_1 i_2 \cdots i_\ell$ and $i_{\ell+1} i_{\ell+2} \cdots i_k$. The outer iteration works with the latter index group whereas the inner iteration works with the former group and executes the classical Yates's algorithm. The pseudocode is as follows:

(1) For each $i_{\ell+1}, i_{\ell+2}, \ldots, i_k \in [t]$ do the following:
   (a) For each $j_1, j_2, \ldots, j_\ell \in [s]$, set
   $$x^{(\ell)}_{j_1 j_2 \cdots j_\ell} \leftarrow 0\,.$$
   (b) For each $j \in D$, set
   $$x^{(\ell)}_{j_1 j_2 \cdots j_\ell} \leftarrow x^{(\ell)}_{j_1 j_2 \cdots j_\ell} + \alpha_{i_{\ell+1} j_{\ell+1}} \alpha_{i_{\ell+2} j_{\ell+2}} \cdots \alpha_{i_k j_k} x_j\,.$$
   (c) Use (classical) Yates's algorithm to compute $u^{(\ell)} \leftarrow A^{\otimes \ell} x^{(\ell)}$.
   (d) For each $i_1, i_2, \ldots, i_\ell \in [t]$, output $u^{(\ell)}_{i_1 i_2 \cdots i_\ell}$ as the value $y_i = y_{i_1 i_2 \cdots i_k}$.

Now take $\ell = \lceil \log_t |D| \rceil$ to achieve the desired output in parts.

Resource consumption. First, observe that the $t^{k-\ell}$ iterations of the outer loop are independent of each other and can be executed in parallel. Each iteration of the outer loop requires $O((t^{\ell+1} + s^{\ell+1})\ell + |D|)$ operations and working space for $O(t^\ell + s^\ell + |D|)$ field elements.

Recalling that $\ell = \lceil (\log |D|)/(\log t) \rceil$ and that $t \geq s$, we observe that each iteration of the outer loop takes $O(|D| t^2 (\log |D|)/(\log t))$ operations and working space for $O(t|D|)$ field elements.

In total the algorithm takes $O\big(t^{k+2} (\log |D|)/(\log t)\big)$ operations to deliver its output in $\Omega(t^{k-1}/|D|)$ independent parts, each consisting of $O(t|D|)$ field elements.

3.3. **Polynomial extension.** Let us now develop a *polynomial extension* of the split/sparse algorithm in the previous section. In particular, we replace the outer loop with a polynomial indeterminate z. That is, in place of iterating the outer loop, we perform repeated *evaluations* at points $z = z_0 \in \mathbb{F}$, with the intuition that values[16] $z_0 \in [t^{k-\ell}]$ deliver precisely the output that is produced by the outer loop of the split/sparse algorithm during iteration $(i_{\ell+1}, i_{\ell+2}, \ldots, i_k) \in [t^{k-\ell}]$.

Accordingly, we replace references to the matrix A for the *outer* indices $i_{\ell+1}, i_{\ell+2}, \ldots, i_k \in [t]$ with $|D|$ polynomials $\alpha_j(z)$, one polynomial for each $j \in D$.

[16]We assume specific $t^{k-\ell}$ values in $\mathbb{F}$ have been identified with elements of $[t^{k-\ell}]$. For example, such an identification is immediate if $\mathbb{F}$ is a finite field (of large enough) prime order.

We present the pseudocode first in abstract form (with an indeterminate z), with the understanding that concrete executions of the algorithm carry out an *evaluation* at a given point $z = z_0 \in \mathbb{F}$. We continue to assume that $t \geq s$. The abstract pseudocode for the polynomial extension is as follows:

(1) For each $j_1, j_2, \ldots, j_\ell \in [s]$, set
$$x^{(\ell)}_{j_1 j_2 \cdots j_\ell}(z) \leftarrow 0\,.$$
(2) For each $j \in D$, set
$$x^{(\ell)}_{j_1 j_2 \cdots j_\ell}(z) \leftarrow x^{(\ell)}_{j_1 j_2 \cdots j_\ell}(z) + \alpha_j(z) x_j\,.$$
(3) Use (classical) Yates's algorithm to compute $u^{(\ell)}(z) \leftarrow A^{\otimes \ell} x^{(\ell)}(z)$.
(4) For each $i_1, i_2, \ldots, i_\ell \in [t]$, output $u^{(\ell)}_{i_1 i_2 \cdots i_\ell}(z)$.

We observe that $u^{(\ell)}_{i_1 i_2 \ldots i_\ell}(z)$ is a polynomial whose degree is at most the maximum of the degrees of the polynomials $\alpha_j(z)$ for $j \in D$. Let us proceed to define these polynomials.

Towards this end, for $i \in [t^{k-\ell}]$, let us write

$$\Phi_i(z) = \prod_{\substack{w=1 \\ w \neq i}}^{t^{k-\ell}} \frac{z-w}{i-w} \tag{6}$$

for the Lagrange interpolation polynomial of degree at most $t^{k-\ell}-1$ that is 1 at i and 0 elsewhere in $[t^{k-\ell}]$. For each $j \in D$ the univariate polynomial

$$\alpha_j(z) = \sum_{i \in [t^{k-\ell}]} \alpha_{i_{\ell+1} j_{\ell+1}} \alpha_{i_{\ell+2} j_{\ell+2}} \cdots \alpha_{i_k j_k} \Phi_i(z) \tag{7}$$

now interpolate over the desired coefficients. Indeed, for all $i \in [t^{k-\ell}]$ and $j \in D$ we observe that

$$\alpha_j(i) = \sum_{i \in [t^{k-\ell}]} \alpha_{i_{\ell+1} j_{\ell+1}} \alpha_{i_{\ell+2} j_{\ell+2}} \cdots \alpha_{i_k j_k}\,.$$

To turn the abstract pseudocode into a concrete evaluation algorithm, we need a fast algorithm for evaluating the $|D|$ polynomials $\alpha_j(z)$ at a specific point $z = z_0$. Suppose we are given $z_0 \in \mathbb{F}$ as input. We need to compute $\alpha_j(z_0) \in \mathbb{F}$ for each $j \in D$. Recalling (7), we have, for all $j \in D$,

$$\alpha_j(z_0) = \sum_{i \in [t^{k-\ell}]} \alpha_{i_{\ell+1} j_{\ell+1}} \alpha_{i_{\ell+2} j_{\ell+2}} \cdots \alpha_{i_k j_k} \Phi_i(z_0)\,. \tag{8}$$

We observe that (8) in fact describes a matrix-vector multiplication: we multiply the $|D| \times t^{k-\ell}$ matrix with entries $\alpha_{i_{\ell+1} j_{\ell+1}} \alpha_{i_{\ell+2} j_{\ell+2}} \cdots \alpha_{i_k j_k}$ with the $t^{k-\ell} \times 1$ vector with entries $\Phi_i(z_0)$ to obtain as a result the $|D| \times 1$ vector with entries $\alpha_j(z_0)$.

Let us now observe that the matrix entries $\alpha_{i_{\ell+1} j_{\ell+1}} \alpha_{i_{\ell+2} j_{\ell+2}} \cdots \alpha_{i_k j_k}$ only require the last $k-\ell$ components of each index $j \in D$. Accordingly, it

suffices to consider a $s^{k-\ell} \times t^{k-\ell}$ matrix and then use $j \in D$ to *index* the computed values to obtain $\alpha_j(z_0)$ for each $j \in D$. In particular, we can use classical Yates's algorithm to compute the value $\alpha_{j_{\ell+1}j_{\ell+2}\cdots j_k}(z_0)$ for each $(j_{\ell+1}, j_{\ell+2}, \cdots, j_k) \in [s]^{k-\ell}$. Recalling that we assume $t \geq s$, this matrix–vector multiplication takes $O(t^{k-\ell+1}(k-\ell) + |D|)$ operations and working space for $O(t^{k-\ell} + |D|)$ field elements.

To initialize Yates's algorithm, we still require the $t^{k-\ell} \times 1$ vector, so let us describe how to compute the values $\Phi_i(z_0)$ for $i \in [t^{k-\ell}]$. If $z_0 \in [t^{k-\ell}]$, the computation is immediate (insert a 1 to position z_0 and fill the rest of the vector with 0s), so let us assume that $z_0 \notin [t^{k-\ell}]$. First, precompute the values $F_0, F_1, \ldots, F_{t^{k-\ell}-1}$ with the recurrence

$$F_0 = 1, \quad F_\ell = \ell \cdot F_{\ell-1}\,.$$

Second, precompute the product

$$\Gamma(z_0) = \prod_{\ell=1}^{t^{k-\ell}} (z_0 - \ell)\,.$$

Finally, observe from (6) that we can compute, for each $i \in [t^{k-\ell}]$,

$$\Phi_i(z_0) = \frac{1}{(-1)^{t^{k-\ell}-i} F_{i-1} F_{t^{k-\ell}-i}} \cdot \frac{\Gamma(z_0)}{(z_0 - i)}\,.$$

This initialization takes in total $O(t^{k-\ell})$ operations and space.

For a $z_0 \in \mathbb{F}$, a vector x, and the matrix A given as input, the polynomial extension algorithm thus runs in $O(|D|t^2(\log|D|)/(\log d) + t^{k-\ell+1}(k-\ell))$ operations and working space for $O(t|D| + t^{k-\ell})$ field elements. The algorithm produces as output the values $u^{(\ell)}_{i_1 i_2 \cdots i_\ell}(z_0)$ for each $i_1, i_2, \ldots, i_\ell \in [t]$.

Viewed as a polynomial in z, each polynomial $u^{(\ell)}_{i_1 i_2 \cdots i_\ell}(z)$ for $i_1, i_2, \ldots, i_\ell \in [t]$ has degree at most $t^{k-\ell} - 1$. Furtermore, for all $i \in [t^k]$ we have $y_i = u^{(\ell)}_{i_1 i_2 \cdots i_\ell}(i_{\ell+1} i_{\ell+1} \cdots i_k)$. Here $y = A^{\otimes k} x$. That is, the polynomial extension algorithm recovers the output of the split/sparse algorithm.

4. Fast and space-efficient evaluation of the $\binom{6}{2}$-linear form

This section studies a multilinear form that integrates over a function or constraint system with only pairwise interactions between variables by decomposing the interactions into $\binom{6}{2} = 15$ parts. The advantage of such a decomposition is that one can use fast matrix multiplication to arrive at a nontrivial arithmetic circuit for fast integration. An example application is counting small subgraphs of a graph, which we will illustrate in more detail in the next section.

4.1. **The problem and the Nešetřil–Poljak formula.** Let χ be an[17] $N \times N$ matrix over a commutative ring $\mathscr{R}$. We seek to compute the $\binom{6}{2}$-linear form

$$X_{\binom{6}{2}} = \sum_{a,b,c,d,e,f} \chi_{ab}\chi_{ac}\chi_{ad}\chi_{ae}\chi_{af}\chi_{bc}\chi_{bd}\chi_{be}\chi_{bf}\chi_{cd}\chi_{ce}\chi_{cf}\chi_{de}\chi_{df}\chi_{ef} \,. \tag{9}$$

A direct evaluation of (9) takes $O(N^6)$ operations and $O(N^2)$ space. Nešetřil and Poljak [25] observe that we can precompute the three $N^2 \times N^2$ matrices

$$\begin{aligned} U_{ab,cd} &= \chi_{ab}\chi_{ac}\chi_{ad}\chi_{bc}\chi_{bd} \,, \\ S_{ab,ef} &= \chi_{ae}\chi_{af}\chi_{be}\chi_{bf}\chi_{ef} \,, \\ T_{cd,ef} &= \chi_{cd}\chi_{ce}\chi_{cf}\chi_{de}\chi_{df} \,, \end{aligned}$$

and then use fast matrix multiplication to compute

$$X_{\binom{6}{2}} = \sum_{a,b,c,d} U_{ab,cd}V_{ab,cd}, \qquad V_{ab,cd} = \sum_{e,f} S_{ab,ef}T_{cd,ef} \,.$$

This takes $O(N^{2\omega+\epsilon})$ arithmetic operations and $O(N^4)$ space for any constant $\epsilon > 0$.

4.2. **A new summation formula and its complexity.** Our main contribution in this section is a new circuit design that matches the Nešetřil–Poljak design in the asymptotic number of arithmetic operations but reduces the space complexity from $O(N^4)$ to $O(N^2)$. In particular, the space complexity is linear in the input size. Furthermore, the new design is easily parallelizable and extendable to a proof polynomial, as we will witness in the next section.

The new design works with the following explicit decomposition of the matrix multiplication tensor. For $r = 1, 2, \ldots, R$, let $\alpha_{de}(r), \beta_{ef}(r), \gamma_{df}(r)$ be ring elements that satisfy the polynomial identity

$$\sum_{d,e,f} u_{de}v_{ef}w_{df} = \sum_{r=1}^{R} \Bigg(\sum_{d,e'} \alpha_{de'}(r)u_{de'} \Bigg) \Bigg(\sum_{e,f'} \beta_{ef'}(r)v_{ef'} \Bigg) \Bigg(\sum_{d',f} \gamma_{d'f}(r)w_{d'f} \Bigg) \,. \tag{10}$$

We can assume that $R = O(N^{\omega+\epsilon/2})$ for an arbitrary constant $\epsilon > 0$, where ω is the limiting exponent for the tensor rank of the matrix multiplication tensor $\langle n, n, n\rangle$ in $\mathscr{R}$ [10, 22].

[17] The present formulation admits an immediate generalization to $\binom{6}{2}$ distinct $N \times N$ matrices, but to keep the notation concise, we work with a single matrix χ.

The new design is as follows. For each $r = 1, 2, \ldots, R$, compute, using fast matrix multiplication,

$$
\begin{aligned}
A_{ab}(r) &= \sum_{d} \chi_{ad}\chi_{bd}H_{ad}(r)\,, \qquad & H_{ad}(r) &= \sum_{e'} \alpha_{de'}(r)\chi_{ae'}\chi_{de'}\,,\\
B_{bc}(r) &= \sum_{e} \chi_{be}\chi_{ce}K_{be}(r)\,, \qquad & K_{be}(r) &= \sum_{f'} \beta_{ef'}(r)\chi_{bf'}\chi_{ef'}\,,\\
C_{ac}(r) &= \sum_{f} \chi_{af}\chi_{cf}L_{cf}(r)\,, \qquad & L_{cf}(r) &= \sum_{d'} \gamma_{d'f}(r)\chi_{cd'}\chi_{d'f}\,.
\end{aligned} \tag{11}
$$

Finally, compute, again using fast matrix multiplication,

$$
P(r) = \sum_{a,b} \chi_{ab}A_{ab}(r)Q_{ab}(r)\,, \qquad Q_{ab}(r) = \sum_{c} \chi_{ac}\chi_{bc}B_{bc}(r)C_{ac}(r)\,. \tag{12}
$$

Each term $P(r)$ takes $O(N^{\omega+\epsilon/2})$ operations and $O(N^2)$ space to compute.

Theorem 13. $X_{\binom{6}{2}} = \sum_{r=1}^{R} P(r)$.

Proof. Expanding and applying (10) for each a, b, c in turn, we have

$$
\begin{aligned}
\sum_{r=1}^{R} P(r) &= \sum_{r=1}^{R} \sum_{a,b,c} \chi_{ab}\chi_{ac}\chi_{bc}A_{ab}(r)B_{bc}(r)C_{ac}(r)\\
&= \sum_{r=1}^{R} \sum_{a,b,c} \chi_{ab}\chi_{ac}\chi_{bc} \sum_{d,e'} \alpha_{de'}(r)\chi_{ad}\chi_{ae'}\chi_{bd}\chi_{de'}\\
&\qquad \sum_{e,f'} \beta_{ef'}(r)\chi_{be}\chi_{bf'}\chi_{ce}\chi_{ef'}\\
&\qquad \sum_{d',f,} \gamma_{d'f}(r)\chi_{af}\chi_{cd'}\chi_{cf}\chi_{d'f}\\
&= \sum_{a,b,c} \chi_{ab}\chi_{ac}\chi_{bc} \sum_{r=1}^{R} \left(\sum_{d,e'} \alpha_{de'}(r)\chi_{ad}\chi_{ae'}\chi_{bd}\chi_{de'} \right)\\
&\qquad \left(\sum_{e,f'} \beta_{ef'}(r)\chi_{be}\chi_{bf'}\chi_{ce}\chi_{ef'} \right)\\
&\qquad \left(\sum_{d',f} \gamma_{d'f}(r)\chi_{af}\chi_{cd'}\chi_{cf}\chi_{d'f} \right)\\
&= \sum_{a,b,c} \chi_{ab}\chi_{ac}\chi_{bc} \sum_{d,e,f} \chi_{ad}\chi_{ae}\chi_{af}\chi_{bd}\chi_{be}\chi_{bf}\chi_{cd}\chi_{ce}\chi_{cf}\chi_{de}\chi_{df}\chi_{ef}\,.
\end{aligned}
$$

□

Thus, we can compute (9) from χ in $O(N^{2\omega+\epsilon})$ operations and $O(N^2)$ space. Furthermore, easy parallelization using up to $O(N^{\omega+\epsilon/2})$ compute nodes is now possible since the values $P(r)$ can be computed independently of each other.

5. A Camelot algorithm for counting small cliques

This section proves Theorem 1 and Theorem 2. In particular, we demonstrate how to reduce counting k-cliques into an evaluation of the $\binom{6}{2}$-linear form, and then develop a proof polynomial $P(x)$ for the $\binom{6}{2}$-linear form, together with a fast evaluation algorithm that computes $P(x_0) \pmod q$ for given $x_0 \in \mathbb{Z}_q$.

5.1. **Reduction to the $\binom{6}{2}$-linear form.** Suppose that 6 divides k. We start with a routine reduction from counting k-cliques in a given n-vertex graph G to the task of evaluating the $\binom{6}{2}$-linear form for a particular matrix χ derived from G.

Let χ be an $N \times N$ integer matrix with $N = \binom{n}{k/6} \leq n^{k/6}$ and entries defined for all $A, B \in \binom{V(G)}{k/6}$ by

$$\chi_{AB} = \left[A \cup B \text{ is a clique in } G \text{ and } A \cap B = \emptyset\right].$$

That is, the entries of χ indicate all pairs of $k/6$-cliques in G that together form a $k/3$-clique.

The $\binom{6}{2}$-linear form $X_{\binom{6}{2}}$ with input χ now counts each k-clique in G exactly $\binom{k}{k/6, k/6, k/6, k/6, k/6, k/6}$ times. Recalling that $N \leq n^{k/6}$, Theorem 2 now follows by applying the algorithm in §4.2 to compute $X_{\binom{6}{2}}$.

5.2. **A proof polynomial for the $\binom{6}{2}$-linear form.** We now turn the counting algorithm in the previous section into a Camelot algorithm. First, we present a proof polynomial $P(x)$ in the Camelot framework that enables us to compute and verify the value $X_{\binom{6}{2}}$ from the common input χ.

We define a univariate polynomial $P(x)$ so that we can recover $X_{\binom{6}{2}}$ via Theorem 13 using evaluations of $P(x)$ at integer points $x = 1, 2, \ldots, R$. We proceed by extending the coefficients $\alpha_{de}(r), \beta_{ef}(r), \gamma_{df}(r)$ into interpolation polynomials and then extending (11) and (12) into polynomials.

Towards this end, for $r = 1, 2, \ldots, R$, let us write

$$\Lambda_r(x) = \prod_{\substack{j=1 \\ j \neq r}}^{R} \frac{x - j}{r - j} \tag{13}$$

for the Lagrange interpolation polynomials of degree at most R for the points $1, 2, \ldots, R$. The univariate polynomials

$$\begin{aligned} \alpha_{de}(x) &= \sum_{r=1}^{R} \alpha_{de}(r)\Lambda_r(x)\,, \\ \beta_{ef}(x) &= \sum_{r=1}^{R} \beta_{ef}(r)\Lambda_r(x)\,, \\ \gamma_{df}(x) &= \sum_{r=1}^{R} \gamma_{df}(r)\Lambda_r(x) \end{aligned} \tag{14}$$

now interpolate over the coefficients $\alpha_{de}(r), \beta_{ef}(r), \gamma_{df}(r)$ when $x = 1, 2, \ldots, R$. From (14), (11), and (12) it is immediate that the univariate polynomials

$$\begin{aligned} A_{ab}(x) &= \sum_{d} \chi_{ad}\chi_{bd}H_{ad}(x)\,, \qquad H_{ad}(x) = \sum_{e'} \alpha_{de'}(x)\chi_{ae'}\chi_{de'}\,, \\ B_{bc}(x) &= \sum_{e} \chi_{be}\chi_{ce}K_{be}(x)\,, \qquad K_{be}(x) = \sum_{f'} \beta_{ef'}(x)\chi_{bf'}\chi_{ef'}\,, \\ C_{ac}(x) &= \sum_{f} \chi_{af}\chi_{cf}L_{cf}(x)\,, \qquad L_{cf}(x) = \sum_{d'} \gamma_{d'f}(x)\chi_{cd'}\chi_{d'f}\,, \end{aligned} \tag{15}$$

and

$$P(x) = \sum_{a,b} \chi_{ab}A_{ab}(r)Q_{ab}(x)\,, \quad Q_{ab}(x) = \sum_{c} \chi_{ac}\chi_{bc}B_{bc}(x)C_{ac}(x) \tag{16}$$

have degree at most $3R$ and the evaluations of $P(x)$ at $x = 1, 2, \ldots, R$ satisfy Theorem 13. Thus, $P(x)$ is a proof polynomial that we can employ in the Camelot framework. For the modulus q we can select one or more primes $q \geq 3R + 1$ to enable interpolation and reconstruction of $X_{\binom{6}{2}}$ over the integers. In particular we can assume $q = O(R)$. Indeed, since $q \geq N^2$ and χ has entries in $\{0, 1\}$, we can recover the integer $0 \leq X_{\binom{6}{2}} = \sum_{r=1}^{R} P(r) \leq N^6$ from evaluations $\sum_{r=1}^{R} P(r) \pmod q$ for at most $O(1)$ distinct primes $q \geq 3R + 1$ using the Chinese Remainder Theorem. This gives proof size $\tilde{O}(R)$.

5.3. **The evaluation algorithm.** To complete the Camelot algorithm it remains to describe how each node evaluates, modulo q, the proof polynomial $P(x)$ at a given point. Let $x_0 \in \mathbb{Z}_q$ be given. We seek to compute $P(x_0) \pmod q$ in time $O(N^{\omega+\epsilon})$.

Let us first assume that we have available the values $\alpha_{de}(x_0), \beta_{ef}(x_0), \gamma_{df}(x_0)$ for all d, e, f. (This is a nontrivial assumption that we will justify in what follows.) Then, using fast matrix multiplication to evaluate (15) and (16) at $x = x_0$, we obtain $P(x_0)$ in $O(N^{\omega+\epsilon})$ operations and $O(N^2)$ space. Arithmetic operations on scalars modulo $q = O(R)$ can be implemented in $\tilde{O}(1)$

time and space, which is subsumed in the choice of $\epsilon > 0$. Thus, we are done assuming we can justify our earlier assumption.

To compute the values $\alpha_{de}(x_0), \beta_{ef}(x_0), \gamma_{df}(x_0)$ for all d, e, f, we proceed to take a look at the detailed structure of the coefficients $\alpha_{de}(r), \beta_{ef}(r), \gamma_{df}(r)$. We will focus on the coefficients $\alpha_{de}(r)$, the cases for $\beta_{ef}(r)$ and $\gamma_{df}(r)$ are symmetric.

By the properties of tensor rank and the fact that matrix multiplication tensors are closed under taking of Kronecker products, for any constant $\epsilon > 0$ we can assume without loss of generality that $R = O(N^{\omega+\epsilon})$ and that there exist positive integer constants R_0, N_0 and a positive integer t such that $N = N_0^t$ and $R = R_0^t$. In particular, we may identify the indices d, e and the index r with t-digit integers in base N_0 and R_0, respectively, so that if we write d_j, e_j and r_j for the jth digit of d, e and r, respectively, the coefficient $\alpha_{de}(r)$ has the Kronecker product form

$$\alpha_{de}(r) = \prod_{j=1}^{t} \alpha^{(0)}_{d_j e_j}(r_j)\,, \tag{17}$$

where $\alpha^{(0)}$ is a matrix of size $N_0^2 \times R_0$ with integer entries. Thus, we may view (17) as an integer matrix of size $N^2 \times R$ which has been obtained as the t-fold Kronecker power of the constant-sized-and-entried matrix $\alpha^{(0)}$.

Let us now describe how to compute the N^2 coefficients $\alpha_{de}(x_0) \pmod q$ for all d, e. (The computations for $\beta_{ef}(x_0)$ and $\gamma_{df}(x_0)$ are symmetric.)

Let us recall from (14) that $\alpha_{de}(x)$ is an interpolating polynomial for the values (17). That is, we have, for all d, e,

$$\alpha_{de}(x_0) = \sum_{r=1}^{R} \alpha_{de}(r)\Lambda_r(x_0) \pmod q\,. \tag{18}$$

We observe that (18) in fact describes a matrix-vector multiplication: we multiply the $N^2 \times R$ matrix with entries $\alpha_{de}(r)$ with the $R \times 1$ vector with entries $\Lambda_r(x_0)$. Because of the Kronecker structure (17), this matrix-vector product can be computed in $O(Rt)$ operations using Yates's algorithm (§3.1).

To initialize Yates's algorithm, we require the $R \times 1$ vector, so let us start by computing the values $\Lambda_r(x_0) \pmod q$ for $r = 1, 2, \ldots, R$. If $x_0 \in \{1, 2, \ldots, R\}$ the computation is immediate (insert a 1 to position x_0 and fill the rest of the vector with 0s), so let us assume that $x_0 \notin \{1, 2, \ldots, R\}$. First, precompute the values $F_0, F_1, \ldots, F_{R-1}$ with the recurrence

$$F_0 = 1, \quad F_j = j \cdot F_{j-1} \pmod q\,.$$

Second, precompute the product

$$\Gamma(x_0) = \prod_{j=1}^{R} (x_0 - j) \pmod q\,.$$

Finally, observe from (13) that we can compute, for each $r = 1, 2, \ldots, R$,

$$\Lambda_r(x_0) = \frac{1}{(-1)^{R-r} F_{r-1} F_{R-r}} \cdot \frac{\Gamma(x_0)}{(x_0 - r)} \pmod q .$$

This initialization takes in total $O(R)$ operations and space.

By choosing a small enough $\epsilon > 0$ to subsume the polylogarithmic terms, the entire evaluation algorithm runs thus in $O(N^{\omega+\epsilon/2})$ time and space. This completes the Camelot framework for computing and verifying $X_{\binom{6}{2}}$ and thus proves Theorem 1.

6. A Camelot algorithm for the number of triangles

This section proves Theorem 3 and Theorem 4 by showing that an algorithm design of Itai and Rodeh [19] for counting triangles in a graph can be transformed into a Camelot algorithm via the polynomial extension of the split/sparse Yates's algorithm (§3.3).

6.1. **The Itai–Rodeh reduction to trace of a product matrix.** To count triangles in an n-vertex graph, it suffices to compute the matrix trace of the product ABC of three $n \times n$ matrices A, B, C with $\{0, 1\}$ entries. Let us furthermore assume that each of these matrices has at most m nonzero entries (that is, the graph has at most m edges, and $A = B = C$ is the adjacency matrix of the graph). In notation, we must compute $\sum_{i,j,k} a_{ij} b_{jk} c_{ki}$.

6.2. **The trace of a triple product of sparse matrices in parallel.** Let us start by proving Theorem 4 and only then develop the proof polynomial and the Camelot algorithm.

Let $\mathbb{F}$ be a field. For $r = 1, 2, \ldots, R$, let $\alpha_{ij}(r), \beta_{jk}(r), \gamma_{ki}(r)$ be field elements that satisfy the polynomial identity

$$\sum_{i,j,k} a_{ij} b_{jk} c_{ki} = \sum_{r=1}^{R} \Big(\sum_{i,j'} \alpha_{ij'}(r) a_{ij'} \Big) \Big(\sum_{j,k'} \beta_{jk'}(r) b_{jk'} \Big) \Big(\sum_{k,i'} \gamma_{ki'}(r) c_{ki'} \Big) . \tag{19}$$

We can assume that $R = O(n^{\omega+\epsilon})$ for an arbitrary constant $\epsilon > 0$, where ω is the limiting exponent for the tensor rank of the matrix multiplication tensor $\langle n, n, n \rangle$ in $\mathbb{F}$ [10, 22].

Analogously to §5.3, let us take a look at the detailed structure of the coefficients $\alpha_{ij}(r), \beta_{jk}(r), \gamma_{ki}(r)$. We will focus on the coefficients $\alpha_{ij}(r)$, the cases for $\beta_{jk}(r)$ and $\gamma_{ki}(r)$ are symmetric.

By the properties of tensor rank and the fact that matrix multiplication tensors are closed under taking of Kronecker products, for any constant $\epsilon > 0$ we can assume without loss of generality that $R = O(n^{\omega+\epsilon})$ and that there exist positive integer constants R_0, n_0 and a positive integer t such that $R = R_0^t$ and $n = n_0^t$. In particular, we may identify the indices d, e and the index r with t-digit integers in base n_0 and R_0, respectively, so that if we

write i_ℓ, j_ℓ and r_ℓ for the ℓth digit of i, j and r, respectively, the coefficient $\alpha_{ij}(r)$ has the Kronecker product form

$$\alpha_{ij}(r) = \prod_{\ell=1}^{t} \alpha^{(0)}_{i_\ell j_\ell}(r_\ell)\,, \tag{20}$$

where $\alpha^{(0)}$ is a matrix of size $n_0^2 \times R_0$ with integer entries. Thus, we may view (20) as an integer matrix of size $n^2 \times R$ which has been obtained as the t-fold Kronecker power of the constant-sized-and-entried matrix $\alpha^{(0)}$.

It follows that we can use the split/sparse Yates algorithm (§3.2) to produce the values

$$A_r = \sum_{i,j'} \alpha_{ij'}(r) a_{ij'}\,, \qquad r = 1, 2, \ldots, R\,,$$

required in (19) in parts of $O(m)$ values each. By the structure of the split/sparse algorithm, these $O(R/m)$ parts can be produced in parallel, independently of each other. By symmetry, the same applies to the values

$$B_r = \sum_{j,k'} \beta_{jk'}(r) b_{jk'}\,, \qquad r = 1, 2, \ldots, R\,,$$

$$C_r = \sum_{k,i'} \gamma_{ki'}(r) c_{ki'}\,, \qquad r = 1, 2, \ldots, R\,.$$

Theorem 4 now follows by working over the rational numbers[18] since we have $\sum_{i,j,k} a_{ij} b_{jk} c_{ki} = \sum_{r=1}^{R} A_r B_r C_r$ by (19).

6.3. **The proof polynomial and evaluation algorithm.** Let us now prove Theorem 3 by extending the algorithm in the previous section into a proof polynomial and its evaluation algorithm. In essence, we replace the split/sparse algorithm (§3.2) with its polynomial extension (§3.3). Let us assume that $\mathbb{F} = \mathbb{Z}_q$ is a finite field of prime order q.

Again it is convenient to restrict to consider the values A_r, the values B_r and C_r are treated symmetrically.

Let m' be the minimum multiple of R_0 at least m. In particular, m' divides R and $m' = \Theta(m)$. The polynomial extension of Yates's algorithm for the coefficients $A_1, A_2, \ldots, A_R$ works with polynomials $A_{r'}(z)$ of degree at most R/m', where $r' = 1, 2, \ldots, m'$.

Let q be a prime at least $3R/m' + 1$. We can find such a prime in time $\tilde{O}(1)$ [2]. For a given input $z_0 \in \mathbb{Z}_q$, the polynomial extension algorithm produces in time $\tilde{O}(m' + R/m')$ the values $A_1(z_0), A_2(z_0), \ldots, A_{m'}(z_0) \pmod q$. By the structure of the algorithm, for $z_0 \in [R/m']$ and $r' = 1, 2, \ldots, m'$ these values constitute exactly the values $A_1, A_2, \ldots, A_R \pmod q$.

[18] Or sufficiently many distinct prime moduli to enable reconstruction using the Chinese Remainder Theorem.

Pursuing the same approach for the values B_r and C_r, we observe that the polynomials $A_{r'}(z), B_{r'}(z), C_{r'}(z)$ have the property that

$$\sum_{z_0 \in [R/m']} \sum_{r'=1}^{m'} A_{r'}(z_0) B_{r'}(z_0) C_{r'}(z_0) = \sum_{r=1}^{R} A_r B_r C_r \,. \tag{21}$$

Let us now define the proof polynomial

$$P(z) = \sum_{r'=1}^{m'} A_{r'}(z) B_{r'}(z) C_{r'}(z) \,,$$

and observe that it has degree at most $3R/m'$.

From (21) and (19) we observe that

$$\sum_{z_0 \in [R/m']} P(z_0) = \sum_{i,j,k} a_{ij} b_{jk} c_{ki} \,.$$

In particular, if we have $P(z) \pmod q$ in coefficient form, we can compute the trace of ABC modulo q in $\tilde{O}(R/m)$ operations. We can recover $P(z) \pmod q$ in coefficient form from at least $3R/m + 1$ evaluations in $\tilde{O}(R/m)$ operations. Each evaluation takes $\tilde{O}(m + R/m)$ operations. Since A, B, C are $\{0,1\}$-valued matrices, the trace of ABC is $O(n^3)$. Up to polylogarithmic factors we may assume that q grows at least as fast as a root function of n. Thus, we can use $O(1)$ distinct primes q and the Chinese Remainder Theorem to recover the trace of ABC from evaluations modulo q. This gives proof size $\tilde{O}(R/m)$. To complete the proof of Theorem 3, select a small enough $\epsilon > 0$ to subsume the polylogarithmic factors.

6.4. **Meeting the Alon–Yuster–Zwick bound.** This section proves Theorem 5. In essence, we observe that the Alon–Yuster–Zwick design [4] admits edge-linear parallel execution if we implement the dense part of the algorithm using the split/sparse variant of Yates's algorithm (§3.2).

Let $\epsilon > 0$ be fixed and, with foresight, let $\Delta = m^{(\omega-1)/(\omega+1)}$. Without loss of generality we may assume that the input graph is connected and hence the number of vertices n satisfies $m \geq n - 1$. In time $\tilde{O}(m)$ we can compute the vertex degrees and partition the vertices into *low* degree vertices of degree at most Δ and *high* degree vertices of degree above Δ. It is immediate that there are at most $2m/\Delta$ high-degree vertices. In time $\tilde{O}(m)$ we can compute the sparse representation of the graph induced by the high-degree vertices. In particular, this induced subgraph has at most m/Δ vertices and at most m edges. Using the split/sparse algorithm on this subgraph as in §6.2 with $R = O\big((m/\Delta)^{\omega+\epsilon}\big) = O\big(m^{2(\omega+\epsilon)/(\omega+1)}\big)$, we can count the triangles consisting only of high-degree vertices in per-node time $\tilde{O}(m)$ using $O\big(m^{(\omega-1+\epsilon)/(\omega+1)}\big)$ parallel nodes.

It remains to count the triangles with at least one low-degree vertex. Label the edge ends incident to each low-degree vertex with unique numbers from $1, 2, \ldots, \lfloor \Delta \rfloor$. Node $u = 1, 2, \ldots, \lfloor \Delta \rfloor$ now proceeds as follows. For each

edge e with end-vertices x and y, consider both ends x and y in turn. (We only describe the process for x, the process for y is symmetric.) If x has low degree, let z be the other end-vertex of the edge whose x-end has been labeled with the label u (if any). If z and y are adjacent (which we can check, e.g., by binary search to a sorted list of edges), we have found a triangle x, y, z. We let this triangle contribute to the count of triangles if both (a) e is the (for example, lexicographic) minimum edge in the triangle x, y, z with at least one low-degree end-vertex, and (b) x is the minimum low-degree vertex in e. It follows that we can count the triangles that have at least one low-degree vertex in per-node time $\tilde{O}(m)$ using $O(\Delta)$ parallel nodes.

7. A proof template for partitioning sum-products

This section supplies a template for proof polynomials that underlies our results for the chromatic and Tutte polynomials. We rely on Kronecker substitution to enable succinct univariate encoding of the partitioning property. (See Kedlaya and Umans [21] for a discussion of Kronecker substitution and its inverse.)

7.1. **The problem.** Let $U = \{1, 2, \ldots, n\}$ be an n-element ground set (or *universe*), and let us write 2^U for the set of all subsets of U. Let us call a function $f : 2^U \to \mathbb{Z}$ a *set function*.

Suppose the common input consists of at least n bits, known to all the nodes. Suppose the input defines a set function f such that for any given $X \subseteq U$ any node can compute the value $f(X)$ in time $O^*(1)$. *Furthermore, we assume that based on the input each node can in $O^*(1)$ time compute an upper bound ϕ such that $|f(X)| \leq \phi$ for all $X \subseteq U$.*

For a positive integer t with $t = O^*(1)$, the nodes would like to compute the value of the *t-part partitioning sum product*

$$\sum_{(X_1, X_2, \ldots, X_t)} f(X_1) f(X_2) \cdots f(X_t)\,, \tag{22}$$

where the sum is over all t-tuples $(X_1, X_2, \ldots, X_t) \in 2^U \times 2^U \times \cdots \times 2^U$ such that

$$X_1 \cup X_2 \cup \cdots X_t = U \text{ and } X_i \cap X_j = \emptyset \text{ holds for all } 1 \leq i < j \leq t. \tag{23}$$

(That is, the t-tuple $(X_1, X_2, \ldots, X_t)$ partitions the universe U into t pairwise disjoint *parts* $X_1, X_2, \ldots, X_t$ so that the ordering of the parts is relevant and zero or more empty parts are permitted.)

Remark. On its own, a node can compute (22) in time $O^*(2^n)$ [7].

7.2. **The proof polynomial.** This section defines the proof polynomial $P(x)$ for partitioning sum-products.

We start by partitioning the universe U into two disjoint parts, $U = E \cup B$, where E is the *explicit* set (whose subsets each node will track explicitly)

and B is, by slight abuse of terminology, a set of *bits* (where the node will rely on advice from the other nodes).

Let us assume that the set B of bits is $B = \{0, 1, \ldots, 2^{|B|-1}\}$. With this assumption, select exactly $|B|$ bits, possibly with repetition. Observe that if your bits sum to $2^{|B|} - 1$, then you have selected each bit exactly once (without repetition).

Let us also recall that there are exactly $\binom{m+k-1}{k-1}$ ways to place m identical balls into k distinct bins. Put otherwise, we can form exactly $\binom{m+k-1}{k-1}$ distinct multisets of size m over a ground set of k elements. Let us write $\binom{B}{k}$ for the set of multisets of size k over B. In particular, $|\binom{B}{|B|}| = \binom{2|B|-1}{|B|-1}$ is precisely the number of ways we can select $|B|$ bits out of B, possibly with repetition. Let us write $\sum M$ for the sum of elements in $M \in \binom{B}{|B|}$. With this notation, our previous observation is equivalent the statement that for all $M \in \binom{B}{|B|}$ we have $\sum M = 2^{|B|} - 1$ if and only if $M = B$.

It will be convenient to present the proof polynomial first over the integers, and only then choose appropriate primes q. The proof polynomial is

$$P(x) = p_0 + p_1 x + p_2 x^2 + \ldots + p_d x^d \tag{24}$$

whose coefficients are defined for all $s = 0, 1, \ldots, d$ with $d = 2^{|B|-1}|B|$ by

$$p_s = \sum_{(X_1, X_2, \ldots, X_t)} f(X_1) f(X_2) \cdots f(X_t)\,, \tag{25}$$

where the sum is over all t-tuples $(X_1, X_2, \ldots, X_t) \in 2^U \times 2^U \times \cdots \times 2^U$ that satisfy the multiset equality

$$X_1 + X_2 + \ldots + X_t = E + M \tag{26}$$

for a multiset $M \in \binom{B}{|B|}$ with $\sum M = s$. (Here we use additive notation for multisets to stress that the element multiplicities must agree for each element in the ground set U.)

Remark. We observe that $p_{2^{|B|-1}}$ agrees with (22). Indeed, (23) holds if and only if (26) holds with $B = M$. Furthermore, $B = M$ if and only if $\sum M = 2^{|B|-1}$. Thus, each node needs to get confidence that the coefficients $p_0, p_1, \ldots, p_d$ have been correctly computed, *over the integers*.

Remark 2. Since $P(x)$ has degree at most $2^{|B|-1}|B|$, it suffices to use primes q with $q = O^*(2^{|B|})$ to enable evaluation and reconstruction of the proof modulo q. Such primes can be found in time $O^*(1)$ [2].

Remark 3. Each node can use the upper bound ϕ for the absolute values of the function f, to compute in time $O^*(1)$ the upper bound $2^{nt+1}\phi^t$ for the absolute values in (25). Assuming that $\log q = \Omega(n)$ and $\log q = \Omega(\phi)$, which will be the case in our instantiations of the template, it suffices to work with $O^*(1)$ distinct primes q to reconstruct the *integer* coefficients $p_0, p_1, \ldots, p_d$. Indeed, recall that $t = O^*(1)$ and that $\log \phi = O^*(1)$ since ϕ is computable from the common input in time $O^*(1)$.

7.3. **A template for the evaluation algorithm.** Let $x_0 \in \{0, 1, \ldots, q-1\}$ be given. We now describe how a node computes $P(x_0) \pmod q$, but will leave one implementation detail unspecified, namely how the function g is computed within the desired computational budget, with the understanding that such algorithms will be provided in subsequent sections when the template is instantiated.

Let w_E, w_B be formal indeterminates. (That is, each node will be computing with polynomials in w_E, w_B with integer coefficients normalized to $\{0, 1, \ldots, q-1\}$ modulo q.) Each node computes the function[19] $g : 2^E \to \mathbb{Z}_q[w_E, w_B]$ defined for all $Y \subseteq E$ by

$$g(Y) = \sum_{\substack{X \subseteq U \\ X \cap \bar{E} \subseteq Y}} f(X) w_E^{|X \cap E|} w_B^{|X \cap B|} \prod_{b \in X \cap B} x_0^b \pmod q . \tag{27}$$

(Recall that $B = \{0, 1, \ldots, 2^{|B|-1}\}$ so the expression x_0^b is well-defined.)

Remark. The time budget for computing the function g is $O^*(2^{|E|})$. This will be justified on a case-by-case basis in subsequent sections.

Assuming the node has computed the function g, it next computes

$$a(w_E, w_B) = \sum_{j_E, j_B} a_{j_E, j_B} w_E^{j_E} w_B^{j_B} = \sum_{Y \subseteq E} (-1)^{|E \setminus Y|} g(Y)^t \pmod q . \tag{28}$$

The node is interested on $a_{|E|,|B|}$, the integer coefficient of the monomial $w_E^{|E|} w_B^{|B|}$ in $a(w)$. Using (28) and the computed g directly, the node can compute $a_{|E|,|B|}$ in time $O^*(2^{|E|})$.

From (27), (28), and standard inclusion–exclusion arguments [7] it follows that

$$a_{|E|,|B|} = \sum_{(X_1, X_2, \ldots, X_t)} f(X_1) f(X_2) \cdots f(X_t) \prod_{j=1}^{t} \prod_{b \in X_j \cap B} x_0^b \pmod q , \tag{29}$$

where the sum is over all t-tuples $(X_1, X_2, \ldots, X_t) \in 2^U \times 2^U \times \cdots \times 2^U$ for which there exists a multiset $M \in \left(\!\binom{B}{|B|}\!\right)$ with

$$X_1 + X_2 + \ldots + X_t = E + M . \tag{30}$$

Furthermore, every multiset $M \in \left(\!\binom{B}{|B|}\!\right)$ satisfies $\sum M \in \{0, 1, \ldots, 2^{|B|-1}|B|\}$.

Comparing (25) and (26) with (29) and (30), we observe that

$$P(x_0) = a_{|E|,|B|} \pmod q .$$

That is, we have evaluated the polynomial $P(x)$ at the chosen point x_0 modulo q.

[19] By "computing the function g" we mean that the node prepares a complete table of values of the function for every possible input $Y \subseteq E$.

7.4. **Running time for evaluation.** Assuming the node keeps to its budget for computing the function g, we observe that $P(x_0) \pmod q$ for a given $x_0 \in \{0, 1, \ldots, q-1\}$ can be computed in time $O^*(2^{|B|} + 2^{|E|})$. This is minimized when $|B| = |E|$. Since $n = |U| = |E| + |B|$, we have $|E| = |B| = n/2$.

8. Counting exact set covers

As a warmup to illustrate the template in §7, let us instantiate it for the problem of counting exact set covers.

8.1. **Counting exact set covers via partitioning sum-products.** Suppose the input a set $\mathcal{F}$ of $O^*(2^{n/2})$ subsets of an n-element universe U. Our task is to compute the number of distinct set partitions of U that we can form by using exactly t sets from $\mathcal{F}$, for a given $1 \leq t \leq n$.[20]

To instantiate the template, let us take the set function $f : 2^U \to \mathbb{Z}$ to be the indicator function for the sets in $\mathcal{F}$. That is, for all $X \subseteq U$ we define

$$f(X) = \begin{cases} 1 & \text{if } X \in \mathcal{F}; \\ 0 & \text{otherwise.} \end{cases} \tag{31}$$

In particular, we can take $\phi = 1$. Then, the value of the partitioning sum-product (22) is exactly $t!$ times our desired solution.

8.2. **Computing the node function.** It remains to complete the template for evaluating the function (27) with f defined by (31) in time $O^*(2^{n/2})$. In particular, we instantiate the template with $|E| = |B| = n/2$.

Remark. Computing (27) within the time budget requires a dedicated algorithm since $\mathcal{F}$ may have $O^*(2^{n/2})$ sets, and hence the sums in (27) (across all choices $Y \subseteq E$) may have up to 2^n terms in total, which is well above our time budget.

The node proceeds as follows. Let q and $x_0 \in \{0, 1, \ldots, q-1\}$ be as in the template. Initialize an array g_0 whose entries are indexed by subsets in 2^E. Each entry $g_0(Z)$ for $Z \subseteq E$ will be a polynomial in the indeterminates w_E, w_B and with coefficients in $\mathbb{Z}_q$. We assume that initially $g_0(Z) = 0$ for all $Z \subseteq E$. Next, we iterate over the $O^*(2^{n/2})$ sets in $\mathcal{F}$. For each $X \in \mathcal{F}$, set $g(X \cap E) \leftarrow g(X \cap E) w_E^{|X \cap E|} w_B^{|X \cap B|} x_0^{\sum X \cap B} \pmod q$. This iteration can be implemented to run in time $O^(2^{n/2})$. After the iteration is complete, use Yates's algorithm [36] on g_0 to obtain in time $O^*(2^{n/2})$ the function g with

$$g(Y) = \sum_{Z \subseteq Y} g_0(Z) = \sum_{\substack{X \in \mathcal{F} \\ X \cap E \subseteq Y}} w_E^{|X \cap E|} w_B^{|X \cap B|} x_0^{\sum X \cap B} \pmod q .$$

Thus, the function (27) can be computed in time $O^*(2^{n/2})$.

[20] To avoid degenerate cases, let us assume that $\mathcal{F}$ does not contain the empty set.

9. The chromatic polynomial

Let us now proceed with a somewhat more intricate application. We instantiate the template in §7 for the chromatic polynomial.

9.1. **The chromatic polynomial via partitioning sum-products.** Let G be an undirected graph with n vertices. The chromatic polynomial $\chi_G(t)$ is a polynomial in t of degree at most n. Thus the values of $\chi_G(t)$ at any $n+1$ points suffice to reconstruct G by interpolation.

The value $\chi_G(t)$ for a positive integer $t = 1, 2, \ldots, n+1$ counts the number of mappings $c : V(G) \to \{1, 2, \ldots, t\}$ such that $c(u) \neq c(v)$ holds for all edges $\{u, v\} \in E(G)$. Such mappings are in a bijective correspondence with the t-tuples $(X_1, X_2, \ldots, X_t)$ that partition $U = V(G)$ into t pairwise disjoint parts $X_1, X_2, \ldots, X_t$ such that each $X_i \subseteq U$ is an independent set of G. Indeed, take $X_k = c^{-1}(k)$ for $k = 1, 2, \ldots, t$ to see the correspondence between preimages of c and the parts.

Thus, $\chi_G(t)$ equals the partitioning sum-product (22) if we choose the set function $f : 2^U \to \mathbb{Z}$ to be the indicator function for independent sets in G. That is, for all $X \subseteq U$ we define

$$f(X) = \begin{cases} 1 & \text{if } X \text{ is independent in } G; \\ 0 & \text{otherwise.} \end{cases} \tag{32}$$

In particular, we can take $\phi = 1$ and instantiate the template for $t = 1, 2, \ldots, n+1$ to recover $\chi_G(t)$ from $\chi_G(1), \chi_G(2), \ldots, \chi_G(n+1)$ by polynomial interpolation.

9.2. **Computing the node function.** Suppose that $t = 1, 2, \ldots, n+1$ is fixed. It remains to complete the template for evaluating the function (27) with f defined by (32) in time $O^*(2^{n/2})$. In particular, we instantiate the template with $|E| = |B| = n/2$.

Remark. Computing (27) within the time budget requires a dedicated algorithm since G may have up to 2^n independent sets, and hence the sum (27) may have up to 2^n terms (for $Y = E$), which is well above our time budget.

The intuition that we pursue in what follows is that an arbitrary independent set $I \subseteq E \cup B$ in G remains independent in G when restricted to the induced subgraphs $G[E]$ and $G[B]$. Furthermore, $I \cap B$ is disjoint from the neighborhood of $I \cap E$ in B *but otherwise has no interactions with* $I \cap E$. These observations enable us to compute (27) in parts across the cut (E, B) with the help of Yates's algorithm [36].

The node proceeds as follows. As per the template, the node first selects a prime q and a uniform random integer $x \in \{0, 1, \ldots, q-1\}$. Then, the node computes the function $\hat{f}_B : 2^B \to \mathbb{Z}_q[w_E, w_B]$ defined for all $X \subseteq B$ by

$$f_B(X) = \begin{cases} \prod_{b \in X} w_B x_0^b & \text{if } X \text{ is independent in } G; \\ 0 & \text{otherwise} \end{cases} \pmod q .$$

This takes time $O^*(2^{|B|})$. Next, using Yates's algorithm [36], the node computes in time $O^*(2^{|B|})$ the function $g_B : 2^B \to \mathbb{Z}_q[w_E, w_B]$ defined for all $Y \subseteq B$ by

$$g_B(Y) = \sum_{X \subseteq Y} f_B(X) \pmod q .$$

For a subset $X \subseteq E$, let us write $\Gamma_{G,B}(X)$ for the set of all vertices $v \in B$ that have at least one neighbor in the set X in G. Next, using the function g_B that the node has computed, it computes the function $\hat{f}_E : 2^E \to \mathbb{Z}_q[w_E, w_B]$ defined for all $X \subseteq E$ by

(33)
$$\hat{f}_E(X) = \begin{cases} w_E^{|X|} g_B\big(B \setminus \Gamma_{G,B}(X)\big) & \text{if } X \text{ is independent in } G; \\ 0 & \text{otherwise} \end{cases} \pmod q .$$

Remark. The term $g_B\big(B \setminus \Gamma_{G,B}(X)\big)$ in (33) extends an independent set $X \subseteq E$ with contributions from independent sets in B in all possible ways that are compatible with X. This aggregation of contributions across the cut (E, B) is the key to staying within the time budget.

Finally, the node computes the function g in time $O^*(2^{|E|})$ using the function $\hat{f}_E$ and Yates's algorithm. Indeed, for all $Y \subseteq E$ by comparing computations (27), and (32), we observe that

$$g(Y) = \sum_{X \subseteq Y} \hat{f}_E(X) \pmod q .$$

10. The Tutte polynomial

We instantiate the template in §7 for the Tutte polynomial.

10.1. **The Tutte polynomial via partitioning sum-products.** Let G be an undirected graph with n vertices and $O^*(1)$ edges, possibly including loops and parallel edges. First, let us recall that task of computing the Tutte polynomial $T_G(x, y)$ can be reduced to the task of evaluating the partition function of the t-state Potts model with integer edge weights.

Towards this end, associate with G the multivariate polynomial

$$Z_G(t, r) = \sum_{F \subseteq E} t^{c(F)} \prod_{e \in F} r_e ,$$

where t is an indeterminate, r is a vector of indeterminates, with one indeterminate r_e for each edge $e \in E = E(G)$, and $c(F)$ denotes the number of connected components in the subgraph of G with vertex set V and edge set F.

The Tutte polynomial $T_G(x, y)$ can be recovered from $Z_G(t, r)$ by observing [30] that

$$T_G(x, y) = (x-1)^{-c(E)} (y-1)^{-|V|} Z_G(t, r) , \tag{34}$$

where $t = (x-1)(y-1)$ and $r_e = y - 1$ for all $e \in E$.

For integer values $t = 1, 2, \ldots$, Fortuin and Kasteleyn [14] observed that

$$Z_G(t, r) = \sum_{\sigma: V \to \{1,2,\ldots,t\}} \prod_{e \in E} \big(1 + r_e[\sigma(e_1) = \sigma(e_2)]\big)\,.$$

(Here we write e_1 and e_2 for the end-vertices of $e \in E$.) Next, assume that $r_e = r$ for all $e \in E$. Define the function $f : 2^U \to \mathbb{Z}$ for all $X \subseteq U = V(G)$ by

$$f(X) = (1 + r)^{|E(G[X])|}\,. \tag{35}$$

For integer values $r = 1, 2, \ldots$ it holds that

$$Z_G(t, r) = \sum_{(X_1, X_2, \ldots, X_t)} f(X_1) f(X_2) \cdots f(X_t)\,. \tag{36}$$

where the sum is over all t-tuples $(X_1, X_2, \ldots, X_t) \in 2^U \times 2^U \times \cdots \times 2^U$ such that

$$X_1 \cup X_2 \cup \cdots X_t = U \text{ and } X_i \cap X_j = \emptyset \text{ holds for all } 1 \le i < j \le t.$$

Thus, to compute the Tutte polynomial of a graph G it suffices to compute the partitioning sum-product (36) with the inner function (35) for $O^*(1)$ integer points (t, r) with $t = O^*(1)$ and $r = O^*(1)$ to enable interpolation of the Tutte polynomial via (34). In particular, we can take $\phi = O^*(1)$.

10.2. **Computing the node function.** Here we assume that $|E| = 2|B|$. Let q be the chosen prime and $x_0 \in \{0, 1, \ldots, q-1\}$. Let $t = O^*(1)$ and $r = O^*(1)$. Let us clean up the node function g somewhat. To enable an application of Yates's algorithm to get the final node function g, each node needs to compute the function $g_0 : 2^E \to \mathbb{Z}_q[w_E, w_B]$ defined for all $Y \subseteq E$ by

$$g_0(Y) = w_E^{|Y|} \sum_{X \subseteq B} w_B^{|X|} x_0^{\sum X} f(X \cup Y)\,, \pmod q$$

where the inner function f is defined by

$$f(X \cup Y) = \prod_{e \in E(G[X \cup Y])} (1 + r) = (1 + r)^{|E(G[X \cup Y])|}\,.$$

The inner function f in particular involves interactions between X and Y across the cut (E, B), which appear less easy to control/aggregate than similar interactions in the case of independent sets. Each node will rely on the following tripartite strategy to compute g_0. Split E into two disjoint subsets E_1, E_2 with $|E_1| = |E_2|$. Split $Y \subseteq E$ accordingly into $Y_1 = Y \cap E_1$ and $Y_2 = Y \cap E_2$. Since the interactions in f are only between at most two out of the three parts E_1, E_2, B, we have the product decomposition

$$f(X \cup Y_1 \cup Y_2) = f_{B,E_1}(X \cup Y_1) f_{B,E_2}(X \cup Y_2) f_{E_1,E_2}(Y_1 \cup Y_2)\,,$$

where

$$f_{B,E_1}(X \cup Y_1) = (1+r)^{|E(G[X,Y_1])|+|E(G[X])|}\,,$$
$$f_{B,E_2}(X \cup Y_2) = (1+r)^{|E(G[X,Y_2])|+|E(G[Y_2])|}\,,$$
$$f_{E_1,E_2}(Y_1 \cup Y_2) = (1+r)^{|E(G[Y_1,Y_2])|+|E(G[Y_1])|}\,.$$

Let us now augment this decomposition by subsuming elements from the node function g_0 to the terms of the decomposition. Let

$$w_E^{|Y_1|+|Y_2|} w_B^{|X|} x_0^{\sum X} f(X \cup Y_1 \cup Y_2) = \hat{f}_{B,E_1}(X \cup Y_1)\hat{f}_{B,E_2}(X \cup Y_2) f_{E_1,E_2}(Y_1 \cup Y_2)\,,$$

where

$$\hat{f}_{B,E_1}(X \cup Y_1) = (1+r)^{|E(G[X,Y_1])|+|E(G[X])|} w_E^{|Y_1|} w_B^{|X|} x_0^{\sum X}\,,$$
$$\hat{f}_{B,E_2}(X \cup Y_2) = (1+r)^{|E(G[X,Y_2])|+|E(G[Y_2])|} w_E^{|Y_2|}\,,$$
$$f_{E_1,E_2}(Y_1 \cup Y_2) = (1+r)^{|E(G[Y_1,Y_2])|+|E(G[Y_1])|}\,.$$

Assuming that $|E_1| = |E_2| = |B|$, each node can in time $O^*(2^{|E|})$ compute each of the functions $\hat{f}_{B,E_1}$, $\hat{f}_{B,E_2}$, and f_{E_1,E_2}.

For all $Y_1 \subseteq E_1$ and $Y_2 \subseteq E_2$ we have

$$\begin{aligned} g_0(Y_1 \cup Y_2) &= f_{E_1,E_2}(Y_1 \cup Y_2) \sum_{X \subseteq B} \hat{f}_{B,E_1}(X \cup Y_1)\hat{f}_{B,E_2}(X \cup Y_2) \\ &= f_{E_1,E_2}(Y_1 \cup Y_2) t_{E_1,E_2}(Y_1 \cup Y_2)\,, \end{aligned} \tag{37}$$

where the function $t_{E_1,E_2} : 2^{E_1 \cup E_2} \to \mathbb{Z}_q[w_E, w_B]$ is defined for all $Y_1 \subseteq E_1$ and $Y_2 \subseteq E_2$ by

$$t_{E_1,E_2}(Y_1 \cup Y_2) = \sum_{X \subseteq B} \hat{f}_{B,E_1}(X \cup Y_1)\hat{f}_{B,E_2}(X \cup Y_2)\,. \tag{38}$$

In particular, t_{B,E_1} is just a matrix product of the matrices $\hat{f}_{B,E_1}$ and $\hat{f}_{B,E_2}$.

Using (37) and (38), for any constant $\epsilon > 0$ each node can compute the function g_0 in time $O^*(2^{|E_1|(\omega+\epsilon)/3}) = O^*(2^{|E|(\omega+\epsilon)/2}) = O^*(2^{(\omega+\epsilon)n/3})$. (Indeed, recall that we assume $|E_1| = |E_2| = |B|$.) Within the same time budget, a node can recover g by taking the zeta tranform of g_0.

Acknowledgments

We thank Ryan Williams for useful discussions and giving us early access to his manuscript describing his batch evaluation framework [35] of which the present "Camelot" framework forms a special case. We also thank Stefan Schneider and Russell Impagliazzo for an inspiring exposition of the nondeterministic exponential time hypothesis and the Carmosino *et al.* [11] results at the Simons Institute in fall 2015.

References

[1] A. Abboud, R. Williams, H. Yu, More applications of the polynomial method to algorithm design, Proc. 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2015, January 4–6, San Diego, CA), SIAM, Philadelphia, PA, 2015, pp. 218–230.

[2] M. Agrawal, N. Kayal, N. Saxena, PRIMES is in P, *Ann. of Math. (2)* 160 (2004) 781–793.

[3] J. Alman, R. Williams, Probabilistic polynomials and Hamming nearest neighbors, Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2015, 17–20 October 2015, Berkeley, CA), IEEE Computer Society, 2015, pp. 136–150.

[4] N. Alon, R. Yuster, U. Zwick, Finding and counting given length cycles, *Algorithmica* 17 (1997) 209–223.

[5] L. Babai, S. Moran, Arthur-Merlin games: a randomized proof system, and a hierarchy of complexity classes, *J. Comput. System Sci.* 36 (1988) 254–276.

[6] A. Björklund, Below all subsets for some permutational counting problems, arXiv preprint 1211.0391, 2012.

[7] A. Björklund, T. Husfeldt, M. Koivisto, Set partitioning via inclusion-exclusion, *SIAM J. Comput.* 39 (2009) 546–563.

[8] A. Björklund, T. Husfeldt, P. Kaski, M. Koivisto, Computing the Tutte polynomial in vertex-exponential time, Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2008, October 25–28, Philadelphia, PA), IEEE Computer Society, 2008, pp. 677–686.

[9] A. Björklund, T. Husfeldt, P. Kaski, M. Koivisto, Covering and packing in linear space, *Inform. Process. Lett.* 111 (2011) 1033–1036.

[10] P. Bürgisser, M. Clausen, M. A. Shokrollahi, *Algebraic Complexity Theory*, Springer, 1997.

[11] M. L. Carmosino, J. Gao, R. Impagliazzo, I. Mihajlin, R. Paturi, S. Schneider, Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility, Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science (ITCS'16, Cambridge, MA, January 14–16, 2016), ACM, 2016, pp. 261–270.

[12] F. Eisenbrand, F. Grandoni, On the complexity of fixed parameter clique and dominating set, *Theoret. Comput. Sci.* 326 (2004) 57–67.

[13] M. Elkin, An improved construction of progression-free sets, *Israel J. Math.* 184 (2011) 93–128.

[14] C. M. Fortuin, P. W. Kasteleyn, On the random-cluster model. I. Introduction and relation to other models, *Physica* 57 (1972) 536–564.

[15] C. Godsil, G. Royle, *Algebraic Graph Theory*, Springer, 2001.

[16] J. von zur Gathen, J. Gerhard, *Modern Computer Algebra*, 3rd ed., Cambridge University Press, 2013.

[17] S. Gao, A new algorithm for decoding Reed–Solomon codes, in: V. K. Bhargava, H. V. Poor, V. Tarokh, S. Yoon, eds., *Communications, Information, and Network Security*, Springer, 2003, pp. 55–68.

[18] A. Grønlund Jørgensen, S. Pettie, Threesomes, degenerates, and love triangles, Proc. 55th IEEE Annual Symposium on Foundations of Computer Science (FOCS 2014, Philadelphia, PA, 18–21 October 2014), IEEE Computer Society, 2014, pp. 621–630.

[19] A. Itai, M. Rodeh, Finding a minimum circuit in a graph, *SIAM J. Comput.* 7 (1978) 413–423.

[20] R. M. Karp, Dynamic programming meets the principle of inclusion and exclusion, *Oper. Res. Lett.* 1 (1981/82) 49–51.

[21] K. S. Kedlaya, C. Umans, Fast polynomial factorization and modular composition, *SIAM J. Comput.* 40 (2011) 1767–1802.

[22] J. M. Landsberg, *Tensors: Geometry and Applications*, American Mathematical Society, 2012.

[23] F. Le Gall, Powers of tensors and fast matrix multiplication, Proc. International Symposium on Symbolic and Algebraic Computation (ISSAC'14, Kobe, Japan, July 23–25, 2014), ACM, 2014, pp. 296–303.

[24] D. E. Muller, Boolean algebras in electric circuit design, *Amer. Math. Monthly* 61 (1954), no. 2, part II, 27–28.

[25] J. Nešetřil, S. Poljak, On the complexity of the subgraph problem, *Comment. Math. Univ. Carolin.* 26 (1985) 415–419.

[26] M. Pătraşcu, Towards polynomial lower bounds for dynamic problems, Proc. 42nd ACM Symposium on Theory of Computing (STOC'10, Cambridge, MA, 5–8 June 2010), ACM, 2010, pp. 603–609.

[27] I. S. Reed, A class of multiple-error-correcting codes and the decoding scheme, *Trans. I.R.E. PGIT-4* (1954) 38–49.

[28] I. S. Reed, G. Solomon, Polynomial codes over certain finite fields, *J. Soc. Indust. Appl. Math.* 8 (1960) 300–304.

[29] H. J. Ryser, *Combinatorial Mathematics*, Wiley, 1963.

[30] A. D. Sokal, The multivariate Tutte polynomial (alias Potts model) for graphs and matroids, in: *Surveys in Combinatorics, 2005*, Cambridge University Press, 2005, pp. 173–226.

[31] D. Tiwari, S. Gupta, G. Gallarno, J. Rogers, D. Maxwell, Reliability lessons learned from GPU experience with the Titan supercomputer at Oak Ridge leadership computing facility, Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'15, Austin, TX, 15–20 November 2015), Article No. 38.

[32] L. G. Valiant, The complexity of computing the permanent, *Theoret. Comput. Sci.* 8 (1979) 189–201.

[33] D. J. A. Welsh, *Complexity: Knots, Colourings and Counting*, Cambridge University Press, 1993.

[34] R. Williams, A new algorithm for optimal 2-constraint satisfaction and its implications, *Theoret. Comput. Sci.* 348 (2005) 357–365.

[35] R. Williams, Strong ETH breaks with Merlin and Arthur: Short non-interactive proofs of batch evaluation, Electronic Colloquium on Computational Complexity Report TR16-002, January 2016.

[36] F. Yates, *The Design and Analysis of Factorial Experiments*, Imperial Bureau of Soil Science, Harpenden, 1937.

Appendix

Appendix A. An inventory of earlier polynomials

These polynomials are essentially due to Williams [35], except for the set cover polynomial F_t, which is implicit in [7], and for the polynomial for Convolution3SUM [26], which is a minor technical extension of the polynomials due to Williams. In particular, we make no claims of originality for the polynomials in this section but rather present the polynomials to (a) demonstrate the versatility of the Camelot framework, and (b) provide an accessible path towards the main results of the paper.

A.1. **A polynomial for counting Boolean orthogonal vectors.** This section proves Theorem 11(1). Let us recall the Boolean orthogonal vectors problem. We are given two matrices, $A = (a_{ij})$ and $B = (b_{ij})$, both of size $n \times t$ and with $\{0, 1\}$ entries. *For each of the rows in A, we want to count the number of rows in B orthogonal to it.* In notation, we want to find, for each $i = 1, 2, \ldots, n$, the count

$$c_i = \sum_{k=1}^{n} \left[\sum_{j=1}^{t} a_{ij} b_{kj} = 0 \right], \qquad 0 \leq c_i \leq n\,.$$

Let us now define the proof polynomial $P(x)$. We will work modulo a prime $q = \tilde{O}(nt)$. Using fast interpolation (see §2.2), in time $\tilde{O}(nt)$ we can find t univariate polynomials $A_1(x), A_2(x), \ldots, A_t(x)$ of degree at most n such that $A_j(i) = a_{ij} \pmod q$ holds for all $i = 1, 2, \ldots, n$ and $j = 1, 2, \ldots, t$. Let

$$A(x) = (A_1(x), A_2(x), \ldots, A_t(x))\,.$$

In time $\tilde{O}(nt)$ we can evaluate a t-variate polynomial $B(z_1, z_2, \ldots, z_t)$ such that

$$B(z_1, z_2, \ldots, z_t) = \sum_{i=1}^{n} \left[\sum_{j=1}^{t} b_{ij} z_j = 0 \right] \pmod q$$

holds for all $(z_1, z_2, \ldots, z_t) \in \{0, 1\}^t$. Indeed, we can evaluate via the formula

$$B(z_1, z_2, \ldots, z_t) = \sum_{i=1}^{n} \prod_{j=1}^{t} (1 - b_{ij} z_j) \pmod q\,. \tag{39}$$

Let us now take

$$P(x) = B(A(x)) \pmod q\,.$$

To evaluate $P(x)$ at a point $x_0 \in \mathbb{Z}_q$, a node first computes the vector $A(x_0) \in \mathbb{Z}_q^t$ and then uses the evaluation formula (39) to obtain $P(x_0)$ in time $\tilde{O}(nt)$. We observe that $P(x)$ has degree $d \leq nt$. Furthermore, for all $i = 1, 2, \ldots, n$ we have

$$P(i) = B(a_{i1}, a_{i2}, \ldots, a_{it}) = c_i \pmod q\,.$$

Since $0 \le c_i \le n$, from $P(i) \pmod q$ we can recover the integer c_i. From a correct proof $p_0, p_1, \ldots, p_d \in \mathbb{Z}_q$ we can by interpolation in time $\tilde{O}(q)$ thus recover the integer counts $c_1, c_2, \ldots, c_n$.

A.2. **A polynomial for #CNFSAT.** This section proves Theorem 8(1). Let the common input be a CNF formula with v variables $x_1, x_2, \ldots, x_v$ and m clauses $C_1, C_2, \ldots, C_m$. For convenience, let us assume that v is even.

In time $O^*(2^{v/2})$ each node can prepare two matrices $A = (a_{ij})$ and $B = (b_{ij})$, both of size $2^{v/2} \times m$, with the following structure.

For $i = 1, 2, \ldots, 2^{v/2}$ and $j = 1, 2, \ldots, m$, the entry $a_{ij} = 1$ if and only if the ith Boolean assignment to variables $x_1, x_2, \ldots, x_{v/2}$ satisfies *no* literal in C_j; otherwise $a_{ij} = 0$. The entry $b_{ij} = 1$ if and only if the ith Boolean assignment to variables $x_{v/2+1}, x_{v/2+2}, \ldots, x_v$ satisfies *no* literal in C_j; otherwise $b_{ij} = 0$.

Observe now that for a Boolean assignment $(i_1, i_2) \in [2^{v/2}]^2$ to the variables $x_1, x_2, \ldots, x_v$ we have $\sum_{j=1}^{m} a_{i_1 j} b_{i_2 j} = 0$ if and only if (i_1, i_2) satisfies all the clauses.

We can thus reduce the task of counting all the satisfying assignments to the task of counting orthogonal pairs of vectors (see §A.1), with $n = 2^{v/2}$ and $t = m$. Since $m = O^*(1)$, we obtain a Camelot algorithm for #CNFSAT that prepares a proof of size $O^*(2^{v/2})$ in time $O^*(2^{v/2})$.

A.3. **A polynomial for the Hamming distance distribution.** Now that we have had some initial exposure to the framework, let us upgrade the Boolean orthogonal vectors setup to produce the entire Hamming distance distribution over the n^2 pairs of vectors, localized to each of the n vectors in one set of vectors. The technical gist is that we can control the roots of a (Lagrange) interpolating polynomial by supplying them via separate indeterminates $w_1, w_2, \ldots, w_t$ to extract desired features from the input. This section proves Theorem 11(2).

We are given as input given two matrices, $A = (a_{ij})$ and $B = (b_{ij})$, both of size $n \times t$ and with $\{0, 1\}$ entries. *For each of the rows in A, we want to count, for each possible Hamming distance h, the number of rows in B with distance h to the row.* In notation, we want to find, *for each* $i = 1, 2, \ldots, n$ and $h = 0, 1, \ldots, t$, the count

$$c_{ih} = \sum_{k=1}^{n} \left[\sum_{j=1}^{t} (1 - a_{ij}) b_{kj} + a_{ij} (1 - b_{kj}) = h \right], \qquad 0 \le c_{ih} \le n\,.$$

Let us now define the proof polynomial $P(x)$. We will work modulo a prime $q = \tilde{O}(nt^2)$. In time $\tilde{O}(nt^2)$ we can find t univariate polynomials $A_1(x), A_2(x), \ldots, A_t(x)$ of degree at most $n(t+1)$ such that $A_j(i(t+1)+h) = a_{ij} \pmod q$ holds for all $i = 1, 2, \ldots, n$, $h = 0, 1, \ldots, t$ and $j = 1, 2, \ldots, t$.

Also in time $\tilde{O}(nt^2)$ we can find t univariate polynomials $H_1(x), H_2(x), \ldots, H_t(x)$ of degree at most $n(t+1)$ such that for all $i = 1, 2, \ldots, n$ and $h = 0, 1, \ldots, t$

it holds that

$$\big\{H_1\big(i(t+1)+h\big), H_2\big(i(t+1)+h\big), \ldots, H_t\big(i(t+1)+h\big)\big\} = \{0,1,\ldots,t\}\setminus\{h\}\,.$$

Remark. Above it does not matter which polynomial H_j attains which of the t values in $\{0,1,\ldots,t\}\setminus\{h\}$, but this precise set of t values must be correctly attained for each i. For example, we can force the polynomials to take values so that for $j = 1,2,\ldots,t$ each polynomial H_j takes the lowest value $\{0,1,\ldots,t\}\setminus\{h\}$ not already attained.

Let

$$I(x) = \big(A_1(x), A_2(x), \ldots, A_t(x), H_1(x), H_2(x), \ldots, H_t(x)\big)\,.$$

In time $\tilde{O}(nt)$ we can evaluate a $(t+t)$-variate polynomial

$$B(z_1, z_2, \ldots, z_t, w_1, w_2, \ldots, w_t)$$

of degree at most t such that for all $(z_1, z_2, \ldots, z_t) \in \{0,1\}^t$ and $h = 0,1,\ldots,t$ we have[21]

$$B(z_1,z_2,\ldots,z_t,\{0,1,\ldots,t\}\setminus\{h\}) = \big(\Pi^t_{\substack{\ell=0\\ \ell\neq h}}(h-\ell)\big)\sum_{i=1}^{n}\left[\sum_{j=1}^{t}(1-z_j)b_{ij} + z_j(1-b_{ij}) = h\right] \pmod q\,.$$

Indeed, we can evaluate via the formula

$$B(z_1,z_2,\ldots,z_t,w_1,w_2,\ldots,w_t) = \sum_{i=1}^{n}\prod_{\ell=1}^{t}\left(\left(\sum_{j=1}^{t}(1-z_j)b_{ij} + z_j(1-b_{ij})\right) - w_\ell\right) \pmod q\,. \tag{40}$$

Let us now take

$$P(x) = B(I(x)) \pmod q\,.$$

To evaluate $P(x)$ at a point $x_0 \in \mathbb{Z}_q$, a node first computes the vector $I(x_0) \in \mathbb{Z}_q^{2t}$ and then uses the evaluation formula (40) to obtain $P(x_0)$ in time $\tilde{O}(nt^2)$. We observe that $P(x)$ has degree $d \leq nt^2$. Furthermore, for all $i = 1,2,\ldots,n$ and $h = 0,1,\ldots,t$ we have

$$\begin{aligned} P\big(i(t+1)+h\big) &= B\big(a_{i1}, a_{i2}, \ldots, a_{it}, \{0,1,\ldots,t\}\setminus\{h\}\big) \\ &= \big(\Pi^t_{\substack{\ell=0\\ \ell\neq h}}(h-\ell)\big)c_{ih} \pmod q\,. \end{aligned}$$

Since $0 \leq c_{ih} \leq n$ and $\Pi^t_{\ell=0,\ell\neq h}(\ell - h)$ is invertible modulo q, from $P\big(i(t+1)+h\big) \pmod q$ we can recover the integer c_{ih}. From a correct proof $p_0, p_1, \ldots, p_d \in \mathbb{Z}_q$ we can by interpolation in time $\tilde{O}(q)$ thus recover the integer counts c_{ih}.

[21] Here we use the somewhat sloppy notation $\{0,1,\ldots,t\}\setminus\{h\}$ to indicate that we assign these t values in arbitrary order to the t variables $w_1, w_2, \ldots, w_t$ of B.

A.4. **A polynomial for Convolution3SUM.** This section proves Theorem 11(3). The aim is to illustrate that we can extend simple Boolean circuits (a t-bit ripple carry adder in our case) into polynomials modulo q and compose with input-interpolating polynomials to arrive at a Camelot algorithm.

The Convolution3SUM problem (see Pătraşcu [26]) asks, given an array $A[1, 2, \ldots, n]$ of t-bit integers as input, whether there exist indices $i_1, i_2 \in [n]$ with $A[i_1] + A[i_2] = A[i_1 + i_2]$.

Without loss of generality we may assume that $i_1, i_2 \in [n/2]$. For $i \in [n/2]$, define $c_i = |\{\ell \in [n/2] : A[i] + A[\ell] = A[i + \ell]\}|$. The sum $\sum_{i=1}^{n/2} c_i$ is the number of solutions to the instance A. Thus, it suffices to define a proof polynomial $P(x)$ from which we can recover $P(i) = c_i$ for all $i \in [n/2]$. We will work modulo $q = \tilde{O}(nt^2)$.

Let us index the t bit positions of the integers in A by $j = 1, 2, \ldots, t$, where $j = 1$ is the least significant bit and $j = t$ is the most significant. In time $\tilde{O}(nt)$ we can find t univariate polynomials $A_1(x), A_2(x), \ldots, A_t(x)$ of degree at most n such that $A_j(i) \pmod q$ is the value of bit j of $A[i]$ for all $i = 1, 2, \ldots, n$ and $j = 1, 2, \ldots, t$.

In time $\tilde{O}(nt)$ we can evaluate a $3t$-variate polynomial T such that

$$T(y_1, y_2, \ldots, y_t, z_1, z_2, \ldots, z_t, w_1, w_2, \ldots, w_t) = [y + z = w]$$

holds for all bit-vectors[22] $(y_1, y_2, \ldots, y_t) \in \{0, 1\}^t$, $(z_1, z_2, \ldots, z_t) \in \{0, 1\}^t$, and $(w_1, w_2, \ldots, w_t) \in \{0, 1\}^t$. Indeed, let

$$\begin{aligned} S(b_1, b_2, b_3) = {} & \\ & (1 - b_1)(1 - b_2)b_3 + (1 - b_1)b_2(1 - b_3) + b_1(1 - b_2)(1 - b_3) + b_1b_2b_3 \\ M(b_1, b_2, b_3) = {} & (1 - b_1)b_2b_3 + b_1(1 - b_2)b_3 + b_1b_2(1 - b_3) + b_1b_2b_3 \end{aligned}$$

be the 3-variate *sum* and *majority* polynomials, respectively, and define the *ripple carry* polynomials by the recurrence

$$c_0 = 0, \qquad c_j = M(y_j, z_j, c_{j-1}) \tag{41}$$

for $j = 1, 2, \ldots, t$.

For given $(y_1, y_2, \ldots, y_t) \in \mathbb{Z}_q^t$, $(z_1, z_2, \ldots, z_t) \in \mathbb{Z}_q^t$, and $(w_1, w_2, \ldots, w_t) \in \mathbb{Z}_q^t$ we can now evaluate T via the formula

$$\begin{aligned} & T(y_1, y_2, \ldots, y_t, z_1, z_2, \ldots, z_t, w_1, w_2, \ldots, w_t) = \\ & \qquad (1 - c_t) \prod_{j=1}^{t} \big((1 - w_j)(1 - S(y_j, z_j, c_{j-1})) + w_j S(y_j, z_j, c_{j-1})\big) . \end{aligned} \tag{42}$$

This takes $O(t^2)$ operations if we use the recurrence (41) to evaluate the carries c_j. We observe that T is a $3t$-variate polynomial of degree $O(t^2)$.

[22] Caveat: Recall that y_1 is the *least* significant bit and y_t is the *most* significant bit by our convention. We write $[y + z = w]$ for the $\{0, 1\}$-valued indicator variable that indicates whether $y + z = w$ holds for the t-bit binary integers y, z, w.

Let us now define the proof polynomial

$$P(x) = \sum_{\ell=1}^{n/2} T(A(x), A(\ell), A(x+\ell)) \pmod q .$$

We observe that $P(x)$ has degree $d = O(nt^2)$. To evaluate $P(x)$ at a point $x_0 \in \mathbb{Z}_q$, a node first computes the vector $A(x_0) \in \mathbb{Z}_q^t$ and the vectors $A(x_0+\ell) \in \mathbb{Z}_q^t$ for each $\ell \in [n/2]$. This takes time $\tilde{O}(nt)$. The node then uses the evaluation formula (42) to obtain $P(x_0)$ in time $\tilde{O}(nt^2)$. Furthermore, for all $i = 1, 2, \ldots, n$ we have

$$P(i) = c_i \pmod q .$$

Since $0 \leq c_i \leq n/2$ and $q \geq n$, we can recover c_i from $P(i) \pmod q$.

A.5. **A polynomial for the permanent.** This polynomial illustrates the use of inclusion–exclusion formulas in designing proof polynomials. We prove Theorem 8(2,3). Our starting point is Ryser's formula [29]

$$\operatorname{per} A = \sum_{S \subseteq [n]} (-1)^{n-|S|} \prod_{i=1}^{n} \sum_{j \in S} a_{ij}$$

for the permanent of an $n \times n$ integer matrix $A = (a_{ij})$.

Let us define the proof polynomial $P(x)$. We work modulo multiple distinct primes q with $q = O^*(2^{n/2})$. Such primes can be found in time $O^*(1)$ [2]. Fix one such prime q. In time $O^*(2^{n/2})$ we can find $n/2$ univariate polynomials $D_1(x), D_2(x), \ldots, D_{n/2}(x)$ of degree at most $2^{n/2}$ such that the vector of values

$$D(x) = (D_1(x), D_2(x), \ldots, D_{n/2}(x)) \tag{43}$$

ranges over all the distinct vectors in $\{0,1\}^{n/2}$ modulo q as $x = 0, 1, \ldots, 2^{n/2}-1$. Define the $n/2$-variate polynomial

$$Q(z_1, z_2, \ldots, z_{n/2}) = \sum_{z_{n/2+1}, z_{n/2+2}, \ldots, z_n \in \{0,1\}} (-1)^n \prod_{j=1}^{n} (1 - 2z_j) \prod_{i=1}^{n} \sum_{j=1}^{n} a_{ij} z_j \pmod q . \tag{44}$$

Observe that Q has degree at most n. Let us now take

$$P(x) = Q(D(x)) .$$

To evaluate $P(x)$ at a point $x_0 \in \mathbb{Z}_q$, a node first computes the vector $D(x_0) \in \mathbb{Z}_q^{n/2}$ and then uses the evaluation formula (44) to obtain $P(x_0)$ in time $O^*(2^{n/2})$. We observe that $P(x)$ has degree $d \leq 2^{n/2} n$. Thus it suffices to work modulo $q = O^*(2^{n/2})$ to enable fast evaluation and error-correction. From a correct proof $p_0, p_1, \ldots, p_d \in \mathbb{Z}_q$ we can in time $O^*(q)$ recover the

values $P(0), P(1), \ldots, P(q-1) \in \mathbb{Z}_q$. However, recovering the permanent is slightly more technical. From Ryser's formula, *over the integers* we have

$$\operatorname{per} A = \sum_{z_1,z_2,\ldots,z_{n/2} \in \{0,1\}} Q(z_1, z_2, \ldots, z_{n/2}) = \sum_{i=0}^{2^{n/2}-1} P(i)\,.$$

Since $|\operatorname{per} A| \leq 2^n (\max_{i,j} |a_{ij}|)^n$, one of our chosen primes q alone may not be sufficient to reconstruct $\operatorname{per} A$ from $\operatorname{per} A \pmod q$. Since $q = O^*(2^{n/2})$, we observe that presenting the proof modulo $O^*(1)$ distinct primes q suffices for reconstruction via the Chinese Remainder Theorem.

A similar approach works for counting the number of Hamiltonian paths in a graph [20]. (We omit the detailed proof.)

A.6. **A polynomial for counting set covers.** This section proves Theorem 9. Let us look at another example how to work with inclusion–exclusion formulas. Besides illustrating the use of the method, this example also serves to motivate why our more elaborate designs to work with larger set families are justified.

The set cover counting problem asks, given a set $\mathcal{F}$ of subsets of $[n]$ and an integer $0 \leq t \leq n$ as input, for the number $c_t(\mathcal{F})$ of t-tuples $(X_1, X_2, \ldots, X_t) \in \mathcal{F}^k$ with $X_1 \cup X_2 \cup \cdots X_t = [n]$. We have the inclusion–exclusion formula [7]

$$c_t(\mathcal{F}) = \sum_{Y \subseteq [n]} (-1)^{n-|Y|} |\{X \in \mathcal{F} : X \subseteq Y\}|^t\,.$$

Let us now define the proof polynomial $P(x)$. Again we will be working with multiple primes q with $q = O^*(2^{n/2})$. Recall the vector of polynomials $D(x)$ defined in (43). Define the $n/2$-variate polynomial

$$F_t(y_1, y_2, \ldots, y_{n/2}) = \sum_{y_{n/2+1}, y_{n/2+2}, \ldots, y_n \in \{0,1\}} (-1)^n \prod_{j=1}^{n} (1 - 2y_j) \Bigg(\sum_{X \in \mathcal{F}} \prod_{j=1}^{n} \big(1 - (1 - y_j)[j \in X]\big) \Bigg)^t \pmod q\,. \tag{45}$$

We observe that the polynomial F_t has degree at most $(1+t)n/2$. Let us now take

$$P(x) = F_t(D(x))\,. \tag{46}$$

To evaluate $P(x)$ at a point $x_0 \in \mathbb{Z}_q$, a node first computes the vector $D(x_0) \in \mathbb{Z}_q^{n/2}$ and then uses the evaluation formula (45) to obtain $P(x_0)$ in time $O^*(2^{n/2})$. We observe that $P(x)$ has degree $d \leq 2^{n/2}(1+t)n/2$. Thus it suffices to work modulo $q = O^*(2^{n/2})$ to enable fast evaluation and error-correction. From a correct proof $p_0, p_1, \ldots, p_d \in \mathbb{Z}_q$ we can in time

$O^*(q)$ recover the values $P(0), P(1), \ldots, P(q-1) \in \mathbb{Z}_q$. From the inclusion–exclusion formula, *over the integers* we have

$$c_t(\mathcal{F}) = \sum_{y_1,y_2,\ldots,y_{n/2}\in\{0,1\}} F_t(y_1, y_2, \ldots, y_{n/2}) = \sum_{i=0}^{2^{n/2}-1} P(i)\,.$$

Since $|c_t(\mathcal{F})| \leq 2^{n(t+1)}$ and $q = O^*(2^{n/2})$, we observe that presenting the proof modulo $O^*(1)$ distinct primes q suffices for reconstruction via the Chinese Remainder Theorem.

Observe that the time to evaluate F_t at any fixed point $(y_1, y_2, \ldots, y_{n/2}) \in \{0,1\}^{n/2}$ using (45) is $O^*(2^{n/2}|\mathcal{F}|)$, so the running time at each node is $O^*(2^{n/2})$ if we assume $|\mathcal{F}|$ is bounded by a polynomial in n.

Remarks. The explicit sum over $X \in \mathcal{F}$ in (45) forces bad scaling when $\mathcal{F}$ is large. Assuming $\mathcal{F}$ is implicitly defined, in certain cases (for example, suppose that $\mathcal{F}$ is the set of independent sets of an n-vertex graph G) we can execute the sum implicitly, as we will do in §9 and §10 for the chromatic and Tutte polynomials. However, it appears that such implicit summation requires also a considerable change in the structure of the proof polynomial $P(x)$. Indeed, contrast (46) with (24). Furthermore, executing the summation in implicit form relies on Yates's algorithm for the chromatic polynomial, and on a combination of Yates's algorithm and fast matrix multiplication for the Tutte polynomial.

Appendix B. Further applications

This section illustrates further applications of our main results.

B.1. **Enumerating variable assignments to 2-constraints by weight.** This section proves Theorem 12 and illustrates a further application of the $\binom{6}{2}$-linear form. Here it should be noted that only our use of the $\binom{6}{2}$-linear form is novel; the algebraic embedding is due to Williams [34].

Let $z_1, z_2, \ldots, z_n$ be variables that take values over an alphabet of σ symbols. For convenience, let us assume that 6 divides n. Partition the variables into 6 sets $Z_1, Z_2, \ldots, Z_n$ of size $n/6$ each.

Let $\varphi_1, \varphi_2, \ldots, \varphi_m$ be constraints of arity 2, that is, whether each constraint φ_j is satisfied can be determined by looking at the values of exactly 2 of the variables. For a constraint φ_j, let us write x_{j_1} and x_{j_2} with $j_1 < j_2$ for these variables.

We say that the constraint φ_j has *type* (s,t) for $1 \leq s < t \leq 6$ if (s,t) is the lexicographically least pair with the property that $x_{j_1}, x_{j_2} \in Z_s \cup Z_t$. Observe in particular that each constraint has a unique type.

For $1 \leq s < t \leq 6$, let a_s and a_t be assignments of values to the variables in Z_s and Z_t. Let us write $f^{(s,t)}(a_s, a_t)$ for the number of constraints of type (s,t) that are satisfied by a_s, a_t.

Let $N = \sigma^{n/6}$. That is, N is the number of distinct assignments to the variables in Z_s, $s = 1, 2, \ldots, 6$. Let w be a polynomial indeterminate and define for each (s,t) with $1 \le s < t \le 6$ the $N \times N$ matrix $\chi^{(s,t)}(w)$ with entries

$$\chi^{(s,t)}_{a_s,a_t}(w) = w^{f^{(s,t)}(a_s,a_t)}.$$

(Here the rows and columns of $\chi^{(s,t)}$ are indexed by the assignments of values to the variables in Z_s and Z_t, respectively.) Observe that constructing these matrices takes $O^*(2^{2n/6})$ time and space.

Now let us study the $\binom{6}{2}$-linear form over the $\binom{6}{2}$ matrices $\chi^{(s,t)}(w)$. Viewing the form as a polynomial in w, we have

$$\begin{aligned} X_{\binom{6}{2}}(w) &= \sum_{a_1,a_2,\ldots,a_6} \prod_{1 \le s < t \le 6} \chi^{(s,t)}_{a_s,a_t}(w) \\ &= \sum_{a_1,a_2,\ldots,a_6} w^{\sum_{1 \le s < t \le 6} f^{(s,t)}(a_s,a_t)}. \end{aligned}$$

That is, $X_{\binom{6}{2}}(w)$ is a polynomial of degree at most m with the property that the integer coefficient of each monomial w^k is precisely the number of assignments $(a_1, a_2, \ldots, a_6)$ to the n variables that satisfy exactly k of the m constraints. Thus, it suffices to construct $X_{\binom{6}{2}}(w)$ over the integers. We will accomplish this by performing evaluations of $X_{\binom{6}{2}}(w)$ at $m+1$ distinct integer points w_0 and then interpolating over the integers. For $w_0 = 0, 1, \ldots, m$ we observe that $0 \le X_{\binom{6}{2}}(w_0) \le \sigma^n m^m$, so it suffices to use the proof polynomial and the evaluation algorithm for $X_{\binom{6}{2}}$ in §5 with $N = \sigma^{n/6}$ and $O^*(1)$ distinct primes q to recover the integer $X_{\binom{6}{2}}(w_0)$ via the Chinese Remainder Theorem.